Alcatel-Lucent

Service Router | Release 12.0 R4

7710 SR OS Quality of Service Guide

93-0086-10-02 Edition 01

Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.
The information presented is subject to change without notice. Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.
Copyright © 2014 Alcatel-Lucent
All Rights Reserved.

Alcatel·Lucent

This document is protected by copyright. Except as specifically permitted herein, no portion of the provided information can be reproduced in any form, or by any means, without prior written permission from Alcatel-Lucent.
Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.
The information presented is subject to change without notice.
Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

Copyright 2014 Alcatel-Lucent. All rights reserved.

TABLE OF CONTENTS

Preface 15

Getting Started
Alcatel-Lucent 7710 SR-Series Services Configuration Process 20

QoS Policies
QoS Overview 23
QoS Policies 24
Service and Network QoS Policies 28
Network QoS Policies 29
Network Queue QoS Policies 31
Queue Parameters 34
Service Ingress QoS Policies 45
Egress Forwarding Class Override 52
Service Egress QoS Policies 52
Policy Parameters 60
Scheduler Policies 62
Forwarding Classes 72
High-Priority Classes 73
Assured Classes 73
Best-Effort Classes 74
Shared Queues 74
ATM Traffic Descriptor Profiles 74
QoS Policy Entities 75
Frequently Used QoS Terms 76
Configuration Notes 80

Network QoS Policies
Overview 82
Network Ingress Tunnel QoS Override 83
For Tunnel Terminated IP Routing Decisions 83
Normal QoS Operation 83
Tunnel Termination QoS Override Operation 84
Enabling and Disabling Tunnel Termination QoS Override 84
DSCP Marking CPU Generated Traffic 85
Default DSCP Mapping Table 88
Basic Configurations 89
Create a Network QoS Policy 89
Applying Network Policies 91
Default Network Policy Values 92
Service Management Tasks 97
Deleting QoS Policies 97
Remove a Policy from the QoS Configuration 98
Copying and Overwriting Network Policies 98
Editing QoS Policies 100

Resource Allocation for Network QoS policy 101
Network QoS Policy Command Reference 103
Configuration Commands 107
Generic Commands 107
Operational Commands 108
Multi-Link Frame Relay Commands 110
Network QoS Policy Commands 113
Network Ingress QoS Policy Commands 115
Network Egress QoS Policy Commands 122
Network Egress QoS Policy Forwarding Class Commands 124
Self-Generated Traffic Commands 131
Show Commands 134

Network Queue QoS Policies
Overview 146
Network Queue Parent Scheduler 146
Basic Configurations 147
Create a Network Queue QoS Policy 147
Applying Network Queue Policies 150
MDAs 150
Ethernet Ports 151
SONET/SDH Ports 152
Default Network Queue Policy Values 153
Service Management Tasks 159
Deleting QoS Policies 159
Remove a Policy from the QoS Configuration 159
Copying and Overwriting QoS Policies 160
Editing QoS Policies 165
Network Queue QoS Policy Command Reference 167
Configuration Commands 169
Generic Commands 169
Operational Commands 170
Network Queue QoS Policy Commands 171
Show Commands 187

Service Egress and Ingress QoS Policies
Overview 190
Egress SAP Forwarding Class and Forwarding Profile Overrides 191
SAP Egress QoS Policy Modifications 191
Hardware Support 191
DEI Egress Remarking 192
DEI in IEEE 802.1ad 193
DEI in IEEE 802.1ah 193
IEEE 802.1ad Use Case 194
IEEE 802.1ah Use Case 195
Egress FC-Based Remarking 195
Implementation Requirements 197
Default Service Egress and Egress Policy Values 199
SAP Egress Policy 199

Default SAP Ingress Policy 200
Basic Configurations 202
Create Service Egress and Ingress QoS Policies 202
Service Egress QoS Policy 204
Service Ingress QoS Policy 206
VID Filters 216
Arbitrary Bit Matching of VID Filters 218
QoS and VID Filters 218
Port Group Configuration Example 219
Applying Service Ingress and Egress Policies 220
Service Management Tasks 223
Deleting QoS Policies 223
Remove a QoS Policy from Service SAP(s) 223
Copying and Overwriting QoS Policies 225
Remove a Policy from the QoS Configuration 226
Editing QoS Policies 226
Service SAP QoS Policy Command Reference 227
Configuration Commands 237
Generic Commands 237
Operational Commands 238
Service Ingress QoS Policy Commands 240
Service Ingress QoS Policy Forwarding Class Commands 285
Service Ingress QoS Policy Entry Commands 300
IP QoS Policy Match Commands 309
Service Ingress MAC QoS Policy Match Commands 316
Service Egress QoS Policy Forwarding Class Commands 321
Service Queue QoS Policy Commands 345
Show Commands 362

Queue Sharing and Redirection

Queue Sharing and Redirection 380
Supported Platforms 380
Queue Group Applications 381
Access SAP Queue Group Applications 381
Network Port Queue Groups for IP Interfaces 382
Pseudowire Shaping for Layer 2 and Layer 3 Services 383
Queue Group Templates and Port Queue Groups 385
Queue Group Templates 385
Port Queue Groups 385
Forwarding Plane Queue Groups 387
Redirection Models 388
Access SAP Forwarding Class Based Redirection 389
Ingress and Egress SAP Forwarding Class Redirection Association Rules 392
Access Queue Group Statistics 396
Network IP Interface Forwarding Class-Based Redirection 397
Egress Network Forwarding Class Redirection Association Rules 397
Egress Network IP Interface Statistics 399
Separate Ingress IPv4 and IPv6 Statistics 399
Ingress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection 400

Egress PW Shaping using Spoke-SDP Forwarding Class-Based Redirection 404
Ingress and Egress PW Statistics 409
Queue Group Behavior on LAG 410
Queue Group Queue Instantiation Per Link 410
Per Link Queue Group Queue Parameters 410
Adding a Queue Group to an Existing LAG 410
Removing a Queue Group from a LAG 410
Adding a Port to a LAG 411
Basic Configurations 413
Configuring an Ingress Queue Group Template 413
Configuring Egress Queue Group Template 414
Applying Ingress Queue Group to SAP Ingress Policy 415
Applying Egress Queue Group to SAP Egress Policy 416
SAP-based Egress Queue Re-direction 417
Configuring Queue Group on Ethernet Access Ingress Port 419
Configuring Overrides 421
Configuring Queue Group on Ethernet Access Egress Port 422
Configuring Queue Group for Network Egress Traffic on Port 423
Configuring Queue Group for Network Ingress Traffic on Forwarding Plane 424
Using Queue Groups to Police Ingress/Egress Traffic on Network Interface 425
Configuring Ingress/Egress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection . . 427
Specifying QoS Policies on Service SAPs 430
Configuration Commands 431
Generic Commands 431
Queue Group Commands 432
Show Commands 456
Monitor Commands 491
QoS Queue Group Template Command Reference 495

QoS Scheduler Policies

Overview 500
Scheduler Policies 500
Egress Port-Based Schedulers 500
Service/Subscriber Egress Port Bandwidth Allocation 502
Service or Subscriber Scheduler Child to Port Scheduler Parent 504
Frame and Packet-Based Bandwidth Allocation 508
Queue Parental Association Scope 510
Service or Subscriber-Level Scheduler Parental Association Scope 510
Network Queue Parent Scheduler 511
Foster Parent Behavior for Orphaned Queues and Schedulers 512
Frame-Based Accounting 513
Operational Modifications 513
Existing Egress Port Based Virtual Scheduling 514
Queue Behavior Modifications for Frame Based Accounting 514
Virtual Scheduler Rate and Queue Rate Parameter Interpretation 514
Configuring Port Scheduler Policies 516
Port Scheduler Structure 516
Special Orphan Queue and Scheduler Behavior 516
Packet to Frame Bandwidth Conversion 516

Aggregate Rate Limits for Directly Attached Queues 518
SAP Egress QoS Policy Queue Parenting 518
Network Queue QoS Policy Queue Parenting 518
Egress Port Scheduler Overrides 519
Applying a Port Scheduler Policy to a Virtual Port 519
Applying Aggregate Rate Limit to a VPORT 521
Applying a Scheduler Policy to a VPORT 521
Weighted Scheduler Group in a Port Scheduler Policy 522
Basic Configurations 524
Create a QoS Scheduler Policy 524
Applying Scheduler Policies 526
Creating a QoS Port Scheduler Policy 530
Configuring Port Parent Parameters 531
Configuring Distributed LAG Rate 533
Service Management Tasks 545
Deleting QoS Policies 545
Removing a QoS Policy from a Customer Multi-Service Site 545
Removing a QoS Policy from SAP(s) 546
Removing a Policy from the QoS Configuration 546
Copying and Overwriting Scheduler Policies 548
Editing QoS Policies 550
QoS Scheduler Policy Command Reference 551
Generic Commands 555
Operational Commands 556
Scheduler Policy Commands 557
Port Scheduler Policy Commands 566
Show Commands 573
Clear Commands 602

Slope QoS Policies

Overview 606
Basic Configurations 607
Create a Slope QoS Policy 607
Applying Slope Policies 609
Default Slope Policy Values 610
Deleting QoS Policies 611
Copying and Overwriting QoS Policies 613
Editing QoS Policies 615
Slope QoS Policy Command Reference 617
Configuration Commands 619
Generic Commands 619
Operational Commands 620
Slope Policy QoS Commands 621
Slope Policy QoS Policy Commands 623
RED Slope Commands 624
Show Commands 627

Shared-Queue QoS Policies

Overview 630

Multipoint Shared Queuing 630
Ingress Queuing Modes of Operation 630
Ingress Service Queuing 631
Basic Configurations 638
Modifying the Default Shared-Queue Policy 638
Applying Shared-Queue Policies 639
Default Shared Queue Policy Values 643
Shared-Queue QoS Policy Command Reference 647
Configuration Commands 649
Generic Commands 649
Shared Queue QoS Commands 650
Show Commands 658

QoS ATM Traffic Descriptor Profiles

Overview 660
ATM Traffic Descriptor Profiles 660
ATM Traffic Management 660
QoS Model for ATM-Based Services 660
ATM Service Categories 661
ATM Traffic Descriptors and QoS Parameters 663
Policing 663
Shaping 663
ATM Queuing and Scheduling 665
Congestion Avoidance 665
Basic Configurations 666
Create an ATM-TD-Profile QoS Policy 666
Applying ATM-TD-Profile Policies 667
Default ATM-TD-Profile Policy Values 669
Service Management Tasks 670
Removing a Profile from the QoS Configuration 670
Copying and Overwriting Profile 670
Editing QoS Policies 671
ATM QoS Policy Command Reference 673
Configuration Commands 675
Generic Commands 675
Operational Commands 676
ATM QoS Policy Commands 677
Show Commands 683

QoS in MC-MLPPP

Overview 692
Basic Configurations 697
Configuring MC-MLPPP 698
QoS in MLFR and FRF.12 Fragmentation 699
QoS in MLFR 699
QoS in FRF.12 End-to-End Fragmentation 701
MLPPP Command Reference 703
Configuration Commands 705
MC-MLPPP Commands 705

Standards and Protocol Support 709

Table of Contents

LIST OF TABLES

Getting Started

Table 1: Configuration Process 20

QoS Policies

Table 2: QoS Policy Types and Descriptions 26
Table 3: Default Network QoS Policy Egress Marking 30
Table 4: Default Network QoS Policy DSCP to Forwarding Class Mappings 30
Table 5: Default Network Queue Policy Definition 32
Table 6: Supported Hardware Rates and CIR/PIR Values for Non-Channelized MDAs 38
Table 7: Supported Hardware Rates and CIR/PIR Values 39
Table 8: Port Rates 40
Table 9: Forwarding Class and Enqueuing Priority Classification Hierarchy Based on Rule Type 47
Table 10: Forwarding Class Classification Based on Rule Type 49
Table 11: Service Ingress QoS Policy IP Match Criteria 50
Table 12: Service Ingress QoS Policy MAC Match Criteria 50
Table 13: MAC Match Ethernet Frame Types 50
Table 14: MAC Match Criteria Frame Type Dependencies 51
Table 15: Default Service Ingress Policy ID 1 Definition 51
Table 16: TAF Impact on Shared Buffer Average Utilization Calculation 58
Table 17: Default Slope Policy Definition 60
Table 18: Default Slope Policy Definition 61
Table 19: Supported Scheduler Policies 64
Table 20: Forwarding Class Scheduler Mapping 66
Table 21: Forwarding Classes 72

Network QoS Policies

Table 22: DSCP/FC Marking 85
Table 23: Network Policy Defaults 92
Table 24: Default DSCP Names to DSCP Value Mapping Table 118
Table 25: Default Class Selector Code Points to DSCP Value Mapping Table 118
Table 26: Show QoS Network Table Output Fields 134
Table 27: Show QoS Network Output Fields 138

Network Queue QoS Policies

Table 28: Network Queue Policy Defaults 153
Table 29: cbs forwarding class defaults 181
Table 30: High-prio-only forwarding class defaults 182
Table 31: Network Queue Labels and Descriptions 187

Service Egress and Ingress QoS Policies

Table 32: Classification to and (Re-)Marking from PHB 197
Table 33: SAP Egress Policy Defaults 199
Table 34: SAP Ingress Policy Defaults 200
Table 35: Out-remark command effect 290
Table 36: IP Protocol Names 305
Table 37: Default FC HSMDA Queue ID Mappings 310

Table 38: Ingress HSMDA Queue Mapping Behavior Based on Forwarding Type 310

Queue Sharing and Redirection

QoS Scheduler Policies

Table 39: Show QoS Scheduler-Policy Output Fields 573
Table 40: Show QoS Schedule-Hierarchy Port Output Fields 579
Table 41: Show QoS Scheduler-Hierarchy SAP Output Fields 583
Table 42: Show QoS Scheduler-Hierarchy Subscriber Output Fields 590
Table 43: Show QoS Scheduler-Stats Customer Output Fields 593
Table 44: Show QoS Scheduler-Stats SAP Output Fields 595
Table 45: Show QoS Scheduler-Stats Subscriber Output Fields 596

Slope QoS Policies

Table 46: Slope Policy Defaults 610
Table 47: Show QoS Slope Policy Output Fields 627

Shared-Queue QoS Policies

Table 48: Shared Queue Policy Defaults 643
Table 49: Show QoS Shared Queue Output Fields 658

QoS ATM Traffic Descriptor Profiles

Table 50: ATM Traffic Descriptors 663
Table 51: ATM-TD-Profile Defaults 669
Table 52: Show Port ATM PVC VPI/VCI Detail Output Fields 687

QoS in MC-MLPPP

Table 53: Default Packet Forwarding Class to MLPPP Class Mapping 692
Table 54: Packet Forwarding Class to MLPPP Class Mapping 692
Table 55: MLPPP Class Queue Threshold Parameters 693
Table 56: MLPPP Class Queue Scheduling Parameters 694
Table 57: MLPPP Ingress QoS Profile: Reassembly Timers (msec) 695
Table 58: Default FR Class Queue Threshold Parameters 699
Table 59: Default FR Class Queue Scheduling Parameters 700
Table 60: Default FR Ingress QoS Profile: Reassembly Timers (msec) 701

LIST OF FIGURES

QoS Policies

Figure 1: 7710 SR Traffic Types 28
Figure 2: Traffic Queuing Model for 3 Queues and 3 Classes 47
Figure 3: Example Configuration — Carrier's Carrier Application 48
Figure 4: Egress Forwarding CLass Override 52
Figure 5: RED Slope Characteristics 56
Figure 6: Virtual Scheduler Internal Bandwidth Allocation 63
Figure 7: Hierarchical Scheduler and Queue Association 68
Figure 8: Scheduler Policy on SAP and Scheduler Hierarchy Creation 69
Figure 9: Scheduler Policy on Customer Site and Scheduler Hierarchy Creation 70

Network QoS Policies

Network Queue QoS Policies

Service Egress and Ingress QoS Policies

Figure 10: DE Bit in the 802.1ad S-TAG 193
Figure 11: DE Aware 802.1ad Access Network 194
Figure 12: DE Aware PBB Topology 195
Figure 13: DEI Processing Ingress into the PE1 SAP 197
Figure 14: Ingress Criteria Classification Directly to Policer 213
Figure 15: VID Filtering Examples 217
Figure 16: Port Groups 219

Queue Sharing and Redirection

QoS Scheduler Policies

Figure 17: Port Level Virtual Scheduler Bandwidth Allocation Based on Priority and CIR 503
Figure 18: Two Scheduler Policy Model for Access Ports 504
Figure 19: Schedulers on SAP or Multi-Service Site Receive Bandwidth From Port Priority Levels 505
Figure 20: Direct Service or Subscriber Association to Port Scheduler Model 507
Figure 21: Port Bandwidth Distribution for Service and Port Scheduler Hierarchies 509
Figure 22: Port Bandwidth Distribution for Direct Queue to Port Scheduler Hierarchy 509
Figure 23: Bandwidth Distribution on Network Port with Port-Based Scheduling 511
Figure 24: Applying a Port Scheduler Policy to a Vport 520

Slope QoS Policies

Shared-Queue QoS Policies

Figure 25: Unicast Service Queue Mapping to Multiple Destination Based Hardware Queues 632
Figure 26: Unicast Service Queuing With Shared Queuing Enabled 634
Figure 27: Multipoint Queue Behavior with Shared Queuing Enabled 635
Figure 28: Multipoint Shared Queuing Using First Pass Unicast Queues 637

QoS ATM Traffic Descriptor Profiles

Figure 29: Hierarchical Scheduling for ATM-Based Services 660

QoS in MC-MLPPP

Figure 30: MLPPP Class Queue Thresholds for In-Profile and Out-of-Profile Packets 693
Figure 31: MLPPP Class Queue Scheduling Scheme 694
Figure 32: FR Class Queue Thresholds for In-Profile and Out-of-Profile Packets 700
Figure 33: FR Class Queue Scheduling for an MLFR Bundle 700
Figure 34: DLC Egress Channel Queue Scheduling 702

Preface

About This Guide

This guide describes the Quality of Service (QoS) provided by the routers and presents examples to configure and implement various protocols and services.

This document is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

Audience

This manual is intended for network administrators who are responsible for configuring the routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Protocols, standards, and services described in this manual include the following:

- CLI concepts
- Quality of Service (QoS) policies and profiles

List of Technical Publications

The 7710 SR documentation set is composed of the following books:

- 7710 SR OS Basic System Configuration Guide
  This guide describes basic system configurations and operations.
- 7710 SR OS System Management Guide
  This guide describes system security and access configurations as well as event logging and accounting logs.
- 7710 SR OS Interface Configuration Guide
  This guide describes card, Media Dependent Adapter (MDA) and port provisioning.
- 7710 SR OS Router Configuration Guide
  This guide describes logical IP routing interfaces and associated attributes such as an IP address, as well as IP and MAC-based filtering, and VRRP and Cflowd.
- 7710 SR OS Routing Protocols Guide
  This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, BGP, and route policies.
- 7710 SR OS MPLS Guide
  This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
- 7710 SR OS Services Guide
  This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, and user services.
- 7710 SR OAM and Diagnostic Guide
- This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- 7710 SR OS Triple Play Guide
  This guide describes Triple Play services and support provided by the 7710 SR and presents examples to configure and implement various protocols and services.
- 7710 SR OS Quality of Service Guide
  This guide describes how to configure Quality of Service (QoS) policy management.

Technical Support

If you purchased a service agreement for your 7710 SR-Series router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, contact your welcome center.

http://www.alcatel-lucent.com/wps/portal/support

Report documentation errors, omissions and comments to:

Documentation.feedback@alcatel-lucent.com

Include document name, version, part number and page(s) affected.

Getting Started

In This Chapter

This chapter provides process flow information to configure Quality of Service (QoS) policies and provision services.

Alcatel-Lucent 7710 SR-Series Services Configuration Process

Table 1 lists the tasks necessary to configure and apply QoS policies. This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 1: Configuration Process

Area	Task	Chapter
Policy configuration	Configuring QoS Policies	
	• Network	Network QoS Policies on page 81
	• Network queue	Network Queue QoS Policies on page 145
	• SAP ingress/SAP egress	Service Egress and Ingress QoS Policies on page 189
	• Scheduler	QoS Scheduler Policies on page 499
	• Slope	Slope QoS Policies on page 605
	• Shared queue	Shared-Queue QoS Policies on page 629
	• ATM traffic descriptor	QoS ATM Traffic Descriptor Profiles on page 659
	• Named pools	Named Pools on page 739
	• Named pools	Named Pools on page 739
Reference	• List of IEEE, IETF, and other proprietary entities	Standards and Protocol Support on page 651

Note: In SR OS 12.0.R4 any function that displays an IPv6 address or prefix changes to reflect rules described in RFC 5952, *A Recommendation for IPv6 Address Text Representation*. Specifically, hexadecimal letters in IPv6 addresses are now represented in lowercase, and the correct compression of all leading zeros is displayed. This changes visible display output compared to previous SR OS releases. Previous SR OS behavior can cause issues with operator scripts that use standard IPv6 address expressions and with libraries that have standard IPv6 parsing as per RFC 5952 rules. See the section on IPv6 Addresses in the Router Configuration Guide for more information.

QoS Policies

In This Chapter

This chapter provides information about Quality of Service (QoS) policy management.

Topics in this chapter include:

- QoS Overview on page 23
- Service and Network QoS Policies on page 28
  - → Network QoS Policies on page 29
  - → Network Queue QoS Policies on page 31
  - → Service Ingress QoS Policies on page 45
  - → Service Egress QoS Policies on page 52
  - → Queue Parameters on page 34
- QoS Policies on page 24
- Scheduler Policies on page 62
  - → Virtual Hierarchical Scheduling on page 64
    - – Single Tier Scheduling on page 65
    - – Hierarchical Scheduler Policies on page 67
- Forwarding Classes on page 72
  - → High-Priority Classes on page 73
  - → Assured Classes on page 73
  - → Best-Effort Classes on page 74
  - → Shared Queues on page 74
- ATM Traffic Descriptor Profiles on page 74
- QoS Policy Entities on page 75

- Configuration Notes on page 80

QoS Overview

Routers are designed with Quality of Service (QoS) mechanisms on both ingress and egress to support multiple customers and multiple services per physical interface. The router has an extensive and flexible capabilities to classify, police, shape and mark traffic.

In the Alcatel-Lucent service router's service model, a service is provisioned on the provider-edge (PE) equipment. Service data is encapsulated and then sent in a service tunnel to the far-end Alcatel-Lucent service router where the service data is delivered.

The operational theory of a service tunnel is that the encapsulation of the data between the two Alcatel Lucent service routers (such as the 7950 XRS, 7750 SR, 7710 SR, 7750 SR MG and 7450 ESS) appear like a Layer 2 path to the service data although it is really traversing an IP or IP/MPLS core. The tunnel from one edge device to the other edge device is provisioned with an encapsulation and the services are mapped to the tunnel that most appropriately supports the service needs.

The router supports eight forwarding classes internally named: Network-Control, High-1, Expedited, High-2, Low-1, Assured, Low-2 and Best-Effort. The forwarding classes are discussed in more detail in Forwarding Classes on page 72.

Router use QoS policies to control how QoS is handled at distinct points in the service delivery model within the device. There are different types of QoS policies that cater to the different QoS needs at each point in the service delivery model. QoS policies are defined in a global context in the router and only take effect when the policy is applied to a relevant entity.

QoS policies are uniquely identified with a policy ID number or name. Policy ID 1 or Policy ID "default" is reserved for the default policy which is used if no policy is explicitly applied.

The QoS policies within the router can be divided into three main types:

- QoS policies are used for classification, defining and queuing attributes and marking.
- Slope policies define default buffer allocations and WRED slope definitions.
- Scheduler policies determine how queues are scheduled.

QoS Policies

Service ingress, service egress, and network QoS policies are defined with a scope of either template or exclusive. Template policies can be applied to multiple SAPs or IP interfaces, whereas, exclusive policies can only be applied to a single entity.

On most systems, the number of configurable SAP ingress and egress QOS policies per system is larger than the maximum number that can be applied per FP. The **tools dump system-resources** output displays the actual number of policies applied on a given FP (noting that the default SAP ingress policy is always applied once for internal use). The **show qos sap-ingress** and **show qos sap-egress** commands can be used to show the number of polices configured.

One service ingress QoS policy and one service egress QoS policy can be applied to a specific SAP. One network QoS policy can be applied to a specific IP interface. A network QoS policy defines both ingress and egress behavior.

Router QoS policies are applied on service ingress, service egress, and network interfaces and define:

Classification rules for how traffic is mapped to queues

- The number of forwarding class queues
- The queue parameters used for policing, shaping, and buffer allocation
- QoS marking/interpretation

There are several types of QoS policies:

- Service ingress
- Service egress
- Network (for ingress and egress)
- Network queue (for ingress and egress)
- ATM traffic descriptor profile
- Scheduler
- Shared queue
- Slope

Service ingress QoS policies are applied to the customer-facing Service Access Points (SAPs) and map traffic to forwarding class queues on ingress. The mapping of traffic to queues can be based on combinations of customer QoS marking (IEEE 802.1p bits, DSCP, and TOS precedence), IP and MAC criteria. The characteristics of the forwarding class queues are defined within the policy as to the number of forwarding class queues for unicast traffic and the queue characteristics. There can be up to eight (8) unicast forwarding class queues in the policy; one for each forwarding class.

A service ingress QoS policy also defines up to three (3) queues per forwarding class to be used for multipoint traffic for multipoint services. In the case of the VPLS, four types of forwarding are supported (which is not to be confused with forwarding classes); unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are flooded to all destinations within the service while the unicast forwarding type is handled in a point-to-point fashion within the service.

Service egress QoS policies are applied to SAPs and map forwarding classes to service egress queues for a service. Up to 8 queues per service can be defined for the 8 forwarding classes. A service egress QoS policy also defines how to remark the forwarding class to IEEE 802.1p bits in the customer traffic.

Network QoS policies are applied to IP interfaces. On ingress, the policy applied to an IP interface maps incoming DSCP and EXP values to forwarding class and profile state for the traffic received from the core network. On egress, the policy maps forwarding class and profile state to DSCP and EXP values for traffic to be transmitted into the core network.

Network queue policies are applied on egress to network ports and channels and on ingress to MDAs. The policies define the forwarding class queue characteristics for these entities.

Service ingress, service egress, and network QoS policies are defined with a scope of either *template* or *exclusive*. Template policies can be applied to multiple SAPs or IP interfaces whereas exclusive policies can only be applied to a single entity.

One service ingress QoS policy and one service egress QoS policy can be applied to a specific SAP. One network QoS policy can be applied to a specific IP interface. A network QoS policy defines both ingress and egress behavior.

If no QoS policy is explicitly applied to a SAP or IP interface, a default QoS policy is applied.

A summary of the major functions performed by the QoS policies is listed in Table 2.

Table 2: QoS Policy Types and Descriptions

Policy Type	Applied at…	Description	Page
Service Ingress	SAP ingress	• Defines up to 32 forwarding class queues and queue parameters for traffic classification. • Defines up to 31 multipoint service queues for broadcast, multicast and destination unknown traffic in multipoint services. • Defines match criteria to map flows to the queues based on combinations of customer QoS (IEEE 802.1p bits, DSCP, TOS Precedence), IP criteria or MAC criteria.	45
Service Egress	SAP egress	• Defines up to 8 forwarding class queues and queue parameters for traffic classification. • Maps one or more forwarding classes to the queues.	52
Network	Router interface	Packets are marked using QoS policies on edge devices. Invoking a QoS policy on a network port allows for the packets that match the policy criteria to be remarked. • Used for classification/marking of MPLS packets. • At ingress, defines MPLS LSP-EXP to FC mapping and 12 meters used by FCs. • At egress, defines FC to MPLS LSP-EXP marking.	29
Network	Ports	• Used for classification/marking of IP packets. • At ingress, defines DSCP or Dot1p to FC mapping and 8 meters. • At egress, defines FC to DSCP or Dot1p marking or both. •	
Network Queue	Network ingress	• Defines forwarding class mappings to network queues and queue characteristics for the queues.	31
Slope	Ports	• Enables or disables the high-slope, low-slope, and non-TCP parameters within the egress or ingress pool.	60
Scheduler	Customer multi-service site Service SAP	• Defines the hierarchy and parameters for each scheduler. • Defined in the context of a tier which is used to place the scheduler within the hierarchy. • Three tiers of virtual schedulers are supported.	62
Shared Queue	SAP ingress	• Shared-queues can be implemented to mitigate the queue consumption on an MDA.	74

Table 2: QoS Policy Types and Descriptions (Continued)

Policy Type	Applied at…	Description	Page
ATM Traffic Descriptor Profile	SAP ingress	• Defines the expected rates and characteristics of traffic. Specified traffic parameters are used for policing ATM cells and for selecting the service category for the per-VC queue.	74
ATM Traffic Descriptor Profile	SAP egress	• Specified traffic parameters are used for scheduling and shaping ATM cells and for selecting the service category for the per-VC queue.	74

Service and Network QoS Policies

The QoS mechanisms within the routers are specialized for the type of traffic on the interface. For customer interfaces, there is service ingress and egress traffic, and for network core interfaces, there is network ingress and network egress traffic (Figure 1).

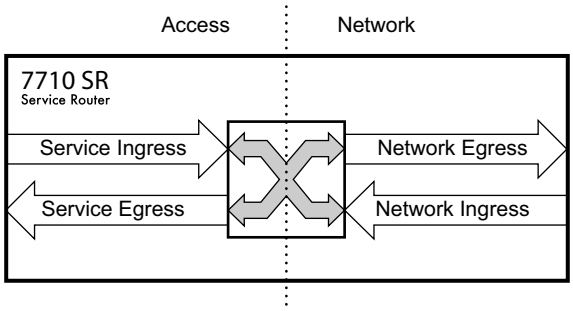

Figure 1: 7710 SR Traffic Types

The router uses QoS policies applied to a SAP for a service or to an network port to define the queuing, queue attributes, and QoS marking/interpretation.

The router supports four types of service and network QoS policies:

- Service ingress QoS policies
- Service egress QoS policies
- Network QoS policies
- Network Queue QoS policies

Network QoS Policies

Network QoS policies define egress QoS marking and ingress QoS interpretation for traffic on core network IP interfaces. The router automatically creates egress queues for each of the forwarding classes on network IP interfaces.

A network QoS policy defines both the ingress and egress handling of QoS on the IP interface. The following functions are defined.

- Ingress
  - → Defines DSCP name mappings to a forwarding classes.
  - → Defines LSP EXP value mappings to forwarding classes.
- Egress
  - → Defines the forwarding class to DSCP value markings.
  - → Defines forwarding class to LSP EXP value markings.
  - → Enables/disables remarking of QoS.

The required elements to be defined in a network QoS policy are:

- A unique network QoS policy ID.
- Egress forwarding class to DSCP value mappings for each forwarding class.
- Egress forwarding class to LSP EXP value mappings for each forwarding class.
- Enabling/disabling of egress QoS remarking.
- A default ingress forwarding class and in-profile/out-of-profile state.

Optional network QoS policy elements include:

- DSCP name to forwarding class and profile state mappings for all DSCP values received.
- LSP EXP value to forwarding class and profile state mappings for all EXP values received.

Network policy ID 1 is reserved as the default network QoS policy. The default policy cannot be deleted or changed.

The default network QoS policy is applied to all network interfaces which do not have another network QoS policy explicitly assigned.

Table 3: Default Network QoS Policy Egress Marking

FC-ID	FC Name	FC Label	DiffServ Name	Egress DSCP Marking		Egress LSP EXP Marking	
				In-Profile Name	Out-of-Profile Name	In-Profile	Out-of-Profile
7	Network Control	nc	NC2	nc2 111000 - 56	nc2 111000 - 56	111 - 7	111 - 7
6	High-1	h1	NC1	nc1 110000 - 48	nc1 110000 - 48	110 - 6	110 - 6
5	Expedited	ef	EF	ef 101110 - 46	ef 101110 - 46	101 - 5	101 - 5
4	High-2	h2	AF4	af41 100010 - 34	af42 100100 - 36	100 - 4	100 - 4
3	Low-1	l1	AF2	af21 010010 - 18	af22 010100 - 20	011 - 3	010 - 2
2	Assured	af	AF1	af11 001010 - 10	af12 001100 - 12	011 - 3	010 - 2
1	Low-2	l2	CS1	cs1 001000 - 8	cs1 001000 - 8	001 - 1	001 - 1
0	Best Effort	be	BE	be 000000 - 0	be 000000 - 0	000 - 0	000 - 0

For network ingress, Table 4 and Table 5 list the default mapping of DSCP name and LSP EXP values to forwarding class and profile state for the default network QoS policy.

Table 4: Default Network QoS Policy DSCP to Forwarding Class Mappings

Ingress DSCP		Forwarding Class			
dscp-name	dscp-value (binary - decimal)	FC ID	Name	Label	Profile State
Default[a]		0	Best-Effort	be	Out
ef	101110 - 46	5	Expedited	ef	In
nc1	110000 - 48	6	High-1	h1	In
nc2	111000 - 56	7	Network Control	nc	In
af11	001010 - 10	2	Assured	af	In

Table 4: Default Network QoS Policy DSCP to Forwarding Class Mappings (Continued)

Ingress DSCP		Forwarding Class			
dscp-name	**dscp-value (binary - decimal)**	**FC ID**	**Name**	**Label**	**Profile State**
af12	001100 - 12	2	Assured	af	Out
af13	001110 - 14	2	Assured	af	Out
af21	010010 - 18	3	Low-1	l1	In
af22	010100 - 20	3	Low-1	l1	Out
af23	010110 - 22	3	Low-1	l1	Out
af31	011010 - 26	3	Low-1	l1	In
af32	011100 - 28	3	Low-1	l1	Out
af33	011110 - 30	3	Low-1	l1	Out
af41	100010 - 34	4	High-2	h2	In
af42	100100 - 36	4	High-2	h2	Out
af43	100110 - 38	4	High-2	h2	Out

Network Queue QoS Policies

Network queue policies define the network forwarding class queue characteristics. Network queue policies are applied on egress on core network ports, channels and on ingress on MDAs. Network queue policies can be configured to use as many queues as needed This means that the number of queues can vary. Not all policies will use eight queues like the default network queue policy.

The queue characteristics that can be configured on a per-forwarding class basis are:

- Committed Buffer Size (CBS) as a percentage of the buffer pool
- Maximum Buffer Size (MBS) as a percentage of the buffer pool
- High Priority Only Buffers as a percentage of MBS
- Peak Information Rate (PIR) as a percentage of egress port bandwidth
- Committed Information Rate (CIR) as a percentage of egress port bandwidth

Network queue policies are identified with a unique policy name which conforms to the standard router alphanumeric naming conventions.

The system default network queue policy is named **default** and cannot be edited or deleted. Table 5 describes the default network queue policy definition.

Table 5: Default Network Queue Policy Definition

Forwarding Class	Queue	Definition
Network-Control (nc)	Queue 8	• PIR = 100% • CIR = 10% • MBS = 25% • CBS = 3% • High-Prio-Only = 10%
High-1 (h1)	Queue 7	• PIR = 100% • CIR = 10% • MBS = 25% • CBS = 3% • High-Prio-Only = 10%
Expedited (ef)	Queue 6	• PIR = 100% • CIR = 100% • MBS = 50% • CBS = 10% • High-Prio-Only = 10%
High-2 (h2)	Queue 5	• PIR = 100% • CIR = 100% • MBS = 50% • CBS = 10% • High-Prio-Only = 10%
Low-1 (l1	Queue 4	• PIR = 100% • CIR = 25% • MBS = 25% • CBS = 3% • High-Prio-Only = 10%
Assured (af)	Queue 3	• PIR = 100% • CIR = 25% • MBS = 50% • CBS = 10% • High-Prio-Only = 10%

Table 5: Default Network Queue Policy Definition (Continued)

Forwarding Class	Queue	Definition (Continued)
Low-2 (l2)	Queue 2	• PIR = 100% • CIR = 25% • MBS = 50% • CBS = 3% • High-Prio-Only = 10%
Best-Effort (be)	Queue 1	• PIR = 100% • CIR = 0% • MBS = 50% • CBS = 1% • High-Prio-Only = 10%

Queue Parameters

This section describes the queue parameters provisioned on access and queues for QoS.

The queue parameters are:

- Queue ID on page 34
- Unicast or Multipoint Queue on page 34
- Queue Hardware Scheduler on page 34
- Committed Information Rate on page 36
- Peak Information Rate on page 37
- Adaptation Rule on page 38
- Committed Burst Size on page 42
- Maximum Burst Size on page 42
- High-Priority Only Buffers on page 42
- Packet Markings on page 43
- Queue-Types on page 44

Queue ID

The queue ID is used to uniquely identify the queue. The queue ID is only unique within the context of the QoS policy within which the queue is defined.

Unicast or Multipoint Queue

Currently, only VPLS services utilize multipoint ingress queues although IES services use multipoint ingress queues for multicast traffic alone when PIM is enabled on the service interface.

Queue Hardware Scheduler

The hardware scheduler for a queue dictates how it will be scheduled relative to other queues at the hardware level. When a queue is defined in a service ingress or service egress QoS policy, it is possible to explicitly define the hardware scheduler to use for the queue when it is applied to a SAP.

Being able to define a hardware scheduler is important as a single queue allows support for multiple forwarding classes. The default behavior is to automatically choose the expedited or non-

expedited nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (nc, ef, h1 or h2), the queue will be treated as an expedited queue by the hardware schedulers. When any non-expedited forwarding classes are mapped to the queue (be, af, l1 or l2), the queue will be treated as best effort by the hardware schedulers.

Committed Information Rate

The committed information rate (CIR) for a queue performs two distinct functions:

1. Profile marking service ingress queues — Service ingress queues mark packets in-profile or out-of-profile based on the queue's CIR. For each packet in a service ingress queue, the CIR is checked with the current transmission rate of the queue. If the current rate is at or below the CIR threshold, the transmitted packet is internally marked in-profile. If the current rate is above the threshold, the transmitted packet is internally marked out-of-profile.
2. Scheduler queue priority metric — The scheduler serving a group of service ingress or egress queues prioritizes individual queues based on their current CIR and PIR states. Queues operating below their CIR are always served before those queues operating at or above their CIR. Queue scheduling is discussed in Virtual Hierarchical Scheduling on page 64.

All router queues support the concept of in-profile and out-of-profile. The network QoS policy applied at network egress determines how or if the profile state is marked in packets transmitted into the service core network. If the profile state is marked in the service core packets, out-of-profile packets are preferentially dropped over in-profile packets at congestion points in the core.

1. When defining the CIR for a queue, the value specified is the administrative CIR for the queue.The router has a number of native rates in hardware that it uses to determine the operational CIR for the queue. The user has some control over how the administrative CIR is converted to an operational CIR should the hardware not support the exact CIR and PIR combination specified. The interpretation of the administrative CIR is discussed below in Adaptation Rule on page 38

Although the router is flexible in how the CIR can be configured, there are conventional ranges for the CIR based on the forwarding class of a queue. A service ingress queue associated with the high-priority class normally has the CIR threshold equal to the PIR rate although the router allows the CIR to be provisioned to any rate below the PIR should this behavior be required. If the service egress queue is associated with a best-effort class, the CIR threshold is normally set to zero; again the setting of this parameter is flexible.

The CIR for a service queue is provisioned on ingress and egress service queues within service ingress QoS policies and service egress QoS policies, respectively.

The CIR for network queues are defined within network queue policies based on the forwarding class. The CIR for the queues for the forwarding class are defined as a percentage of the network interface bandwidth.

Peak Information Rate

The peak information rate (PIR) defines the maximum rate at which packets are allowed to exit the queue. It does not specify the maximum rate at which packets may enter the queue; this is governed by the queue's ability to absorb bursts and is defined by its maximum burst size (MBS).

The actual transmission rate of a service queue depends on more than just its PIR. Each queue is competing for transmission bandwidth with other queues. Each queue's PIR, CIR and the relative importance of the scheduler serving the queue all combine to affect a queue's ability to transmit packets as discussed in Single Tier Scheduling on page 65.

The PIR is provisioned on ingress and egress service queues within service ingress QoS policies and service egress QoS policies, respectively.

The PIR for network queues are defined within network queue policies based on the forwarding class. The PIR for the queues for the forwarding class are defined as a percentage of the network interface bandwidth.

When defining the PIR for a queue, the value specified is the administrative PIR for the queue.The router has a number of native rates in hardware that it uses to determine the operational PIR for the queue. The user has some control over how the administrative PIR is converted to an operational PIR should the hardware not support the exact CIR and PIR values specified. The interpretation of the administrative PIR is discussed below in Adaptation Rule on page 38

Adaptation Rule

The adaptation rule provides the QoS provisioning system with the ability to adapt specific CIR and PIR defined administrative rates to the underlying capabilities of the hardware the queue will be created on to derive the operational rates. The administrative CIR and PIR rates are translated to actual operational rates enforced by the hardware queue. The rule provides a constraint used when the exact rate is not available due to hardware implementation trade-offs.

For the CIR and PIR parameters individually, the system will attempt to find the best operational rate depending on the defined constraint. The supported constraints are:

- Minimum — Find the hardware supported rate that is equal to or higher than the specified rate.
- Maximum — Find the hardware supported rate that is equal to or lesser than the specified rate.
- Closest — Find the hardware supported rate that is closest to the specified rate.

Depending on the hardware upon which the queue is provisioned, the actual operational CIR and PIR settings used by the queue will be dependant on the method the hardware uses to implement and represent the mechanisms that enforce the CIR and PIR rates.

The adaptation rule always assumes that the PIR (shaping parameter) on the queue is the most important rate. When multiple available hardware rates exist for a given CIR and PIR rate pair, the PIR constraint is always evaluated before the CIR.

The router CFMuses a rate step value to define the granularity for both the CIR and PIR rates The adaptation rule controls the method the system uses to choose the rate step based on the administrative rates defined by the **rate** command. The supported CIR and PIR values ranges and increments are summarized in Table 6.

The MDA hardware rate-step values are listed in Table 7 for all MDA egress traffic (except deep channel MDAs).

Table 6: Supported Hardware Rates and CIR/PIR Values for Non-Channelized MDAs

Hardware Rate Steps	Rate Range (Rate Step x 0 to Rate Step x 127 and max)[a]
0.5Gb/sec	0 to 64Gb/sec and ∞
100Mb/sec	0 to 12.7Gb/sec and ∞
50Mb/sec	0 to 6.4Gb/sec and ∞
10Mb/sec	0 to 1.3Gb/sec and ∞
5Mb/sec	0 to 635Mb/sec and ∞
5Mb/sec	0 to 640 MB/sec and ∞

Table 6: Supported Hardware Rates and CIR/PIR Values for Non-Channelized MDAs

Hardware Rate Steps	Rate Range (Rate Step x 0 to Rate Step x 127 and max)[a]
1Mb/sec	0 to 127Mb/sec and ∞
500Kb/sec	0 to 64Mb/sec and ∞
100Kb/sec	0 to 12.7Mb/sec and ∞
50Kb/sec	0 to 6.4Mb/sec and ∞
10Kb/sec	0 to 1.2Mb/sec and ∞
8Kb/sec	0 to 1Mb/sec and ∞
1Kb/sec	0 to 127Kb/sec and ∞

a. 0 is unavailable for PIR

The hardware step-rate values for ingress traffic from all MDAs/CMAs and egress traffic for all CMAs and deep channel MDAs.

The MDA hardware rate-step values are listed below for deep channel MDAs (m1-choc12-sfp, m4-choc3-sfp, and m4-chds3). The table shows supported hardware rates and CIR/PIR values for ingress traffic from all MDAs/CMAs and egress traffic for all CMAs and deep channel MDAs.

Table 7: Supported Hardware Rates and CIR/PIR Values [a]

Hardware Rate Steps	Rate Range
0.5Gb/sec	0 to 64Gb/sec and ∞
100Mb/sec	0 to 12.7Gb/sec and ∞
10Mb/sec	0 to 1.3Gb/sec and ∞ (0 unavailable for PIR)
2Mb/sec	0 to 254Mb/sec and ∞ (0 unavailable for PIR)
1Mb/sec	0 to 127Mb/sec and ∞
500 kb/s	0 to 65 Mb/s and ∞
256 kb/s	0 to 32.5 Mb/s and ∞
128 kb/s	0 to 16.3 Mb/s and ∞
100 kb/s	0 to 12.7 Mb/s and ∞
64 kb/s	0 to 8.1 Mb/s and ∞
50 kb/s	0 to 6.4 Mb/s and ∞
32 kb/s	0 to 4.1 Mb/s and ∞
16 kb/s	0 to 2 Mb/s and ∞

Table 7: Supported Hardware Rates and CIR/PIR Values [a]

Hardware Rate Steps	Rate Range (Continued)
8 kb/s	0 to 1 Mb/s and ∞
4 kb/s	0 to 500 k/s and ∞
1 kb/s	0 to 127 kb/s and ∞

a.Network egress ASAP MDA traffic uses the rate step from Table 6

To illustrate how the adaptation rule constraints **minimum**, **maximum** and **closest** are evaluated in determining the operational CIR or PIR for the router CFM, assume there is a queue where the administrative CIR and PIR values are 401 Mbps and 403 Mbps, respectively. According to Table 6, since the PIR value is given precedence and is in the range of 0 to 635 Mbps, the hardware rate step of 5 Mbps is used.

If the adaptation rule is **minimum**, the operational CIR and PIR values will be 405 Mbps as it is the native hardware rate greater than or equal to the administrative CIR and PIR values.

If the adaptation rule is **maximum**, the operational CIR and PIR values will be 400 Mbps.

If the adaptation rule is **closest**, the operational CIR and PIR values will be 400 Mbps and 405 Mbps, respectively, as those are the closest matches for the administrative values that are even multiples of the 5 Mbps rate step.

The hardware rate step values are for the queue CIR and PIR and have the maximum decrement value of 127.

The port rate is set in a VOQ and hence can use a maximum decrement value of 255.

Table 8: Port Rates

Hardware Rate Steps	Rate Range (Rate Step x 0 to Rate Step x 255 and Max
500Mb/sec	0 to 127.5Gb/sec and ∞
100Mb/sec	0 to 25.5Gb/sec and ∞
50Mb/sec	0 to 12.75Gb/sec and ∞
10Mb/sec	0 to 2.55Gb/sec and ∞
5Mb/sec	0 to 1.275Gb/sec and ∞
1Mb/sec	0 to 255Mb/sec and ∞
500Kb/sec	0 to 127.5Mb/sec and ∞

Table 8: Port Rates (Continued)

Hardware Rate Steps	Rate Range (Rate Step x 0 to Rate Step x 255 and Max
100Kb/sec	0 to 25.5Mb/sec and ∞
50Kb/sec	0 to 12.75Mb/sec and ∞
10Kb/sec	0 to 2.55Mb/sec and ∞
8Kb/sec	0 to 2.04Mb/sec and ∞
1Kb/sec	0 to 255Kb/sec and ∞

Committed Burst Size

The committed burst size (CBS) parameters specify the amount of buffers that can be drawn from the reserved buffer portion of the queue's buffer pool. Once the reserved buffers for a given queue have been used, the queue contends with other queues for additional buffer resources up to the maximum burst size.

The CBS is provisioned on ingress and egress service queues within service ingress QoS policies and service egress QoS policies, respectively. The CBS for a queue is specified in Kbytes.

The CBS for network queues are defined within network queue policies based on the forwarding class. The CBS for the queues for the forwarding class are defined as a percentage of buffer space for the pool.

Maximum Burst Size

The maximum burst size (MBS) parameter specifies the maximum queue depth to which a queue can grow. This parameter ensures that a customer that is massively or continuously over-subscribing the PIR of a queue will not consume all the available buffer resources. For high-priority forwarding class service queues, the MBS can be relatively smaller than the other forwarding class queues because the high-priority service packets are scheduled with priority over other service forwarding classes.

The MBS is provisioned on ingress and egress service queues within service ingress QoS policies and service egress QoS policies, respectively. The MBS for a queue is specified in Kbytes.

The MBS for network queues are defined within network queue policies based on the forwarding class. The MBS for the queues for the forwarding class are defined as a percentage of buffer space for the pool.

High-Priority Only Buffers

High priority (HP)-only buffers are defined on a queue and allow buffers to be reserved for traffic classified as high priority. When the queue depth reaches a specified level, only high-priority traffic can be enqueued. The HP-only reservation for a queue is defined as a percentage of the MBS value.

On service ingress, the HP-only reservation for a queue is defined in the service ingress QoS policy. High priority traffic is specified in the match criteria for the policy.

On service egress, the HP-only reservation for a queue is defined in the service egress QoS policy. Service egress queues are specified by forwarding class. High-priority traffic for a given traffic

class is traffic that has been marked as in-profile either on ingress classification or based on interpretation of the QoS markings.

The HP-only for network queues are defined within network queue policies based on the forwarding class. High-priority traffic for a specific traffic class is marked as in-profile either on ingress classification or based on interpretation of the QoS markings.

Packet Markings

Typically, customer markings placed on packets are not treated as trusted from an in-profile or out-of-profile perspective. This allows the use of the ingress buffering to absorb bursts over PIR from a customer and only perform marking as packets are scheduled out of the queue (as opposed to using a hard policing function that operates on the received rate from the customer). The resulting profile (in or out) based on ingress scheduling into the switch fabric is used by network egress for tunnel marking and egress congestion management.

The high/low priority feature allows a provider to offer a customer the ability to have some packets treated with a higher priority when buffered to the ingress queue. If the queue is configured with a hi-prio-only setting (setting the high priority MBS threshold higher than the queue's low priority MBS threshold) a portion of the ingress queue's allowed buffers are reserved for high priority traffic. An access ingress packet must hit an ingress QoS action in order for the ingress forwarding plane to treat the packet as high priority (the default is low priority).

If the packet's ingress queue is above the low priority MBS, the packet will be discarded unless it has been classified as high priority. The priority of the packet is not retained after the packet is placed into the ingress queue. Once the packet is scheduled out of the ingress queue, the packet will be considered in-profile or out-of-profile based on the dynamic rate of the queue relative to the queue's CIR parameter.

If an ingress queue is not configured with a hi-prio-only parameter, the low priority and high priority MBS thresholds will be the same. There will be no difference in high priority and low priority packet handling. At access ingress, the priority of a packet has no affect on which packets are scheduled first. Only the first buffering decision is affected. At ingress and egress, the current dynamic rate of the queue relative to the queue's CIR does affect the scheduling priority between queues going to the same destination (either the switch fabric tap or egress port). The strict operating priority for queues are (from highest to lowest):

- Expedited queues within the CIR (conform)
- Best Effort queues within the CIR (conform)
- Expedited and Best Effort queues above the CIR (exceed)

For access ingress, the CIR controls both dynamic scheduling priority and marking threshold. At network ingress, the queue's CIR affects the scheduling priority but does not provide a profile

marking function (as the network ingress policy trusts the received marking of the packet based on the network QoS policy).

At egress, the profile of a packet is only important for egress queue buffering decisions and egress marking decisions, not for scheduling priority. The egress queue's CIR will determine the dynamic scheduling priority, but will not affect the packet's ingress determined profile.

Queue Counters

The router maintains counters for queues within the system for granular billing and accounting. Each queue maintains the following counters:

- Counters for packets and octets accepted into the queue
- Counters for packets and octets rejected at the queue
- Counters for packets and octets transmitted in-profile
- Counters for packets and octets transmitted out-of-profile

Queue-Types

The **expedite**, **best-effort** and **auto-expedite** queue types are mutually exclusive to each other. Each defines the method that the system uses to service the queue from a hardware perspective. While parental virtual schedulers can be defined for the queue, they only enforce how the queue interacts for bandwidth with other queues associated with the same scheduler hierarchy. An internal mechanism that provides access rules when the queue is vying for bandwidth with queues in other virtual schedulers is also needed.

Color Aware Profiling (Policing)

The normal handling of SAP ingress access packets applies an in-profile or out-of-profile state to each packet relative to the dynamic rate of the queue as the packet is forwarded towards the egress side of the system. When the queue rate is within or equal to the configured CIR, the packet is considered in-profile. When the queue rate is above the CIR, the packet is considered out-of-profile. (This applies when the packet is scheduled out of the queue, not when the packet is buffered into the queue.) Egress queues use the profile marking of packets to preferentially buffer in-profile packets during congestion events. Once a packet has been marked in-profile or out-of-profile by the ingress access SLA enforcement, the packet is tagged with an in-profile or out-of-profile marking allowing congestion management in subsequent hops towards the packet's ultimate destination. Each hop to the destination must have an ingress table that determines the in-profile or out-of-profile nature of a packet based on its QoS markings.

Color aware profiling adds the ability to selectively treat packets received on a SAP as in-profile or out-of-profile regardless of the queue forwarding rate. This allows a customer or access device to color a packet out-of-profile with the intention of preserving in-profile bandwidth for higher priority packets. The customer or access device may also color the packet in-profile, but this is rarely done as the original packets are usually already marked with the in-profile marking.

Each ingress access forwarding class may have one or multiple sub-class associations for SAP ingress classification purposes. Each sub-class retains the chassis wide behavior defined to the parent class while providing expanded ingress QoS classification actions. Sub-classes are created to provide a match association that enforces actions different than the parent forwarding class. These actions include explicit ingress remarking decisions and color aware functions.

All non-profiled and profiled packets are forwarded through the same ingress access queue to prevent out-of-sequence forwarding. Profiled packets in-profile are counted against the total packets flowing through the queue that are marked in-profile. This reduces the amount of CIR available to non-profiled packets causing fewer to be marked in-profile. Profiled packets out-of-profile are counted against the total packets flowing through the queue that are marked in-profile. This ensures that the amount of non-profiled packets marked out-of-profile is not affected by the profiled out-of-profile packet rate.

Service Ingress QoS Policies

Service ingress QoS policies define ingress service forwarding class queues and map flows to those queues. When a service ingress QoS policy is created by default, it always has two queues defined that cannot be deleted: one for the default unicast traffic and one for the default multipoint traffic. These queues exist within the definition of the policy. The queues only get instantiated in hardware when the policy is applied to a SAP. In the case where the service does not have multipoint traffic, the multipoint queues will not be instantiated.

In the simplest service ingress QoS policy, all traffic is treated as a single flow and mapped to a single queue, and all flooded traffic is treated with a single multipoint queue. The required elements to define a service ingress QoS policy are:

- A unique service ingress QoS policy ID.
- A QoS policy scope of template or exclusive.
- At least one default unicast forwarding class queue. The parameters that can be configured for a queue are discussed in Queue Parameters on page 34.
- At least one multipoint forwarding class queue.

Optional service ingress QoS policy elements include:

- Additional unicast queues up to a total of 32.
- Additional multipoint queues up to 31.
- QoS policy match criteria to map packets to a forwarding class.

To facilitate more forwarding classes, sub-classes are now supported. Each forwarding class can have one or multiple sub-class associations for SAP ingress classification purposes. Each sub-class retains the chassis wide behavior defined to the parent class while providing expanded ingress QoS classification actions.

There can now be up to 64 classes and subclasses combined in a sap-ingress policy. With the extra 56 values, the size of the forwarding class space is more than sufficient to handle the various combinations of actions.

Forwarding class expansion is accomplished through the explicit definition of sub-forwarding classes within the SAP ingress QoS policy. The CLI mechanism that creates forwarding class associations within the SAP ingress policy is also used to create sub-classes. A portion of the sub-class definition directly ties the sub-class to a parent, chassis wide forwarding class. The sub-class is only used as a SAP ingress QoS classification tool, the sub-class association is lost once ingress QoS processing is finished.

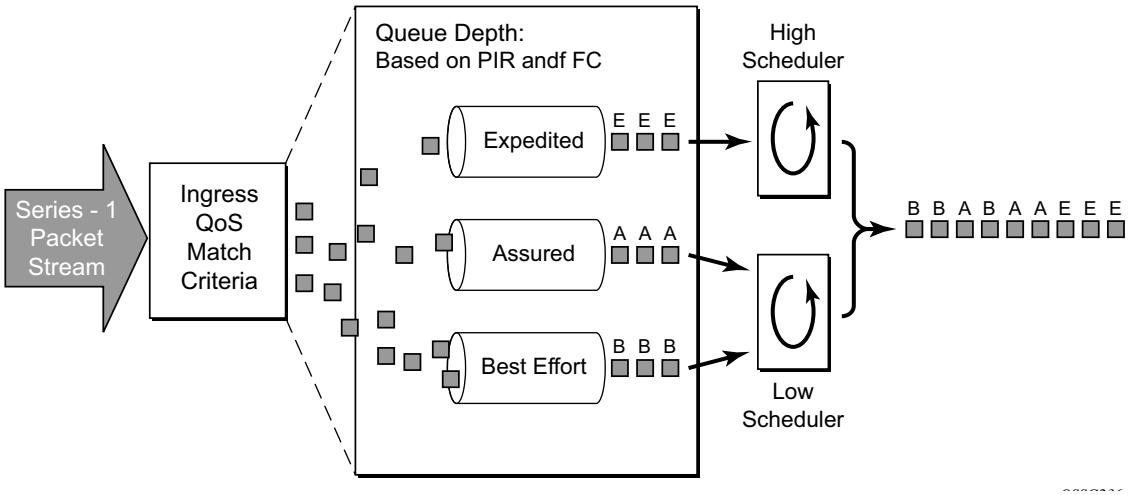

Figure 2: Traffic Queuing Model for 3 Queues and 3 Classes

When configured with this option, the forwarding class and drop priority of incoming traffic will be determined by the mapping result of the EXP bits in the top label. Table 9 displays he new classification hierarchy based on rule type.:

Table 9: Forwarding Class and Enqueuing Priority Classification Hierarchy Based on Rule Type

#	Rule	Forwarding Class	Enqueuing Priority	Comments
1	default-fc	Set the policy's default forwarding class.	Set to policy default	All packets match the default rule.
2	dot1p dot1p-value	Set when an fc-name exists in the policy. Otherwise, preserve from the previous match.	Set when the priority parameter is high or low. Otherwise, preserve from the previous match.	Each dot1p-value must be explicitly defined. Each packet can only match a single dot1p rule.
3	lsp-exp exp-value	Set when an fc-name exists in the policy. Otherwise, preserve from the previous match.	Set when the priority parameter is high or low. Otherwise, preserve from the previous match.	* Each exp-value must be explicitly defined. Each packet can only match a single lsp-exp rule. * This rule can only be applied on Ethernet L2 SAP
4	prec ip-prec-value	Set when an fc-name exists in the policy. Otherwise, preserve from the previous match.	Set when the priority parameter is high or low. Otherwise, preserve from the previous match	Each ip-prec-value value must be explicitly defined. Each packet can only match a single prec rule.

Table 9: Forwarding Class and Enqueuing Priority Classification Hierarchy Based on Rule Type

#	Rule	Forwarding Class	Enqueuing Priority	Comments
5	dscp dscp-name	Set when an fc-name exists in the policy. Otherwise, preserve from the previous match.	Set when the priority parameter is high or low in the entry. Otherwise, preserve from the previous match.	Each dscp-name that defines the DSCP value must be explicitly defined. Each packet can only match a single DSCP rule.
6	IP criteria: Multiple entries per policy Multiple criteria per entry	Set when an fc-name exists in the entry's action. Otherwise, preserve from the previous match.	Set when the priority parameter is high or low in the entry action. Otherwise, preserve from the previous match.	When IP criteria is specified, entries are matched based on ascending order until first match and then processing stops. A packet can only match a single IP criteria entry.
7	MAC criteria: Multiple entries per policy Multiple criteria per entry	Set when an fc-name exists in the entry's action. Otherwise, preserve from the previous match.	Set when the priority parameter is specified as high or low in the entry action. Otherwise, preserve from the previous match.	When MAC criteria is specified, entries are matched based on ascending order until first match and then processing stops. A packet can only match a single MAC criteria entry.

FC Mapping Based on EXP Bits at VLL/VPLS SAP

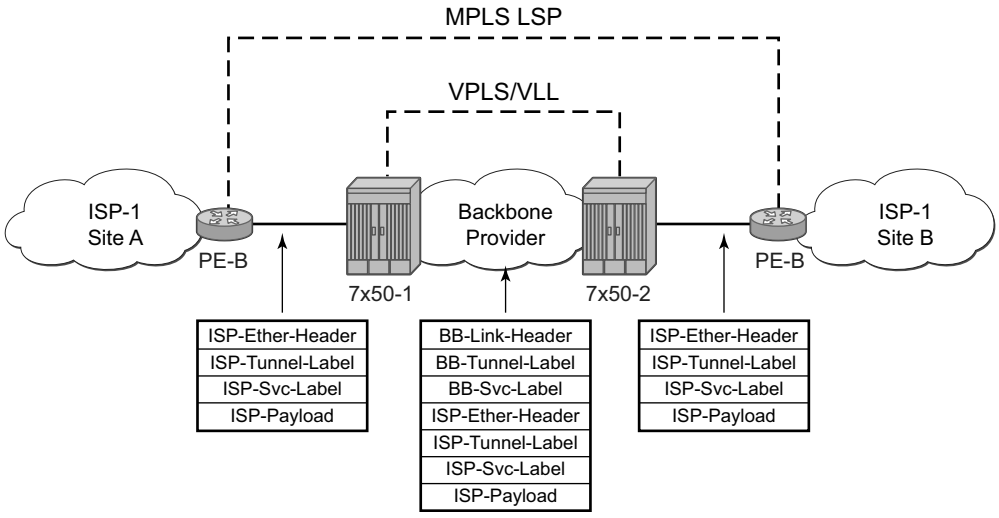

Figure 3: Example Configuration — Carrier's Carrier Application

To accommodate backbone ISPs who want to provide VPLS/VLL to small ISPs as a site-to-site inter-connection service, small ISP routers can connect to a 7x50 Ethernet Layer 2 SAPs. The traffic will be encapsulated in a VLL/VPLS SDP. These small ISP routers are typically PE router. In order to provide appropriate QoS, the 7x50 support a new classification option that based on received MPLS EXP bits.

The **lsp-exp** command is will be supported in sap-ingress qos policy. This option can only be applied on Ethernet Layer 2 SAPs.

Table 10: Forwarding Class Classification Based on Rule Type

#	Rule	Forwarding Class	Comments
1	**default-fc**	Set the policy's default forwarding class.	All packets match the default rule.
2	IP criteria: • Multiple entries per policy • Multiple criteria per entry	Set when an *fc-name* exists in the entry's action. Otherwise, preserve from the previous match.	When IP criteria is specified, entries are matched based on ascending order until first match and then processing stops. A packet can only match a single **IP** criteria entry.
3	MAC criteria: • Multiple entries per policy • Multiple criteria per entry	Set when an *fc-name* exists in the entry's action. Otherwise, preserve from the previous match.	When MAC criteria is specified, entries are matched based on ascending order until first match and then processing stops. A packet can only match a single MAC criteria entry.

The enqueuing priority is specified as part of the classification rule and is set to "high" or "low". The enqueuing priority relates to the forwarding class queue's High-Priority-Only allocation where only packets with a high enqueuing priority are accepted into the queue once the queue's depth reaches the defined threshold. (See High-Priority Only Buffers on page 42.)

The mapping of IEEE 802.1p bits, IP Precedence and DSCP values to forwarding classes is optional as is specifying IP and MAC criteria.

The IP and MAC match criteria can be very basic or quite detailed. IP and MAC match criteria are constructed from policy entries. An entry is identified by a unique, numerical entry ID. A single entry cannot contain more than one match value for each match criteria. Each match entry has a queuing action which specifies: the forwarding class of packets that match the entry.

- The forwarding class of packets that match the entry.
- The enqueuing priority (high or low) for matching packets.

The entries are evaluated in numerical order based on the entry ID from the lowest to highest ID value. The first entry that matches all match criteria has its action performed. Table 11 and Table 12 list the supported IP and MAC match criteria.

Table 11: Service Ingress QoS Policy IP Match Criteria

IP Criteria
• Destination IP address/prefix • Destination port/range • IP fragment • Protocol type (TCP, UDP, etc.) • Source port/range • Source IP address/prefix • DSCP value

Table 12: Service Ingress QoS Policy MAC Match Criteria

MAC Criteria
• IEEE 802.2 LLC SSAP value/mask • IEEE 802.2 LLC DSAP value/mask • IEEE 802.3 LLC SNAP OUI zero or non-zero value • IEEE 802.3 LLC SNAP PID value • IEEE 802.1p value/mask • Source MAC address/mask • Destination MAC address/mask • EtherType value

The MAC match criteria that can be used for an Ethernet frame depends on the frame's format. See Table 13.

Table 13: MAC Match Ethernet Frame Types

Frame Format	Description
802dot3	IEEE 802.3 Ethernet frame. Only the source MAC, destination MAC and IEEE 802.1p value are compared for match criteria.
802dot2-llc	IEEE 802.3 Ethernet frame with an 802.2 LLC header.
802dot2-snap	IEEE 802.2 Ethernet frame with 802.2 SNAP header.
Ethernet-II	Ethernet type II frame where the 802.3 length field is used as an Ethernet type (Etype) value. Etype values are two byte values greater than 0x5FF (1535 decimal).

The 802dot3 frame format matches across all Ethernet frame formats where only the source MAC, destination MAC and IEEE 802.1p value are compared. The other Ethernet frame types match those field values in addition to fields specific to the frame format. Table 14 lists the criteria that can be matched for the various MAC frame types.

Table 14: MAC Match Criteria Frame Type Dependencies

Frame Format	Source MAC	Dest MAC	IEEE 802.1p Value	Etype Value	LLC Header SSAP/DSAP Value/Mask	SNAP-OUI Zero/Non-zero Value	SNAP-PID Value
802dot3	Yes	Yes	Yes	No	No	No	No
802dot2-llc	Yes	Yes	Yes	No	Yes	No	No
802dot2-snap	Yes	Yes	Yes	No	No[a]	Yes	Yes
ethernet-II	Yes	Yes	Yes	Yes	No	No	No

a. When a SNAP header is present, the LLC header is always set to AA-AA

Service ingress QoS policy ID 1 is reserved for the default service ingress policy. The default policy cannot be deleted or changed.

The default service ingress policy is implicitly applied to all SAPs which do not explicitly have another service ingress policy assigned. The characteristics of the default policy are listed in Table 15.

Table 15: Default Service Ingress Policy ID 1 Definition

Characteristic	Item	Definition
Queues	Queue 1	1 (one) queue all unicast traffic: • Forward Class: best-effort (be) • CIR = 0 • PIR = max (line rate) • MBS, CBS and HP Only = default (values derived from applicable policy)
	Queue 11	1 (one) queue for all multipoint traffic: • CIR = 0 • PIR = max (line rate) • MBS, CBS and HP Only = default (values derived from applicable policy)
Flows	Default Forwarding Class	1 (one) flow defined for all traffic: • All traffic mapped to best-effort (be) with a low priority

Egress Forwarding Class Override

Egress forwarding class override provides additional QoS flexibility by allowing the use of a different forwarding class at egress than was used at ingress.

The ingress QoS processing classifies traffic into a forwarding class (or sub-class) and by default the same forwarding class is used for this traffic at the access or network egress. The ingress forwarding class or sub-class can be overridden so that the traffic uses a different forwarding class at the egress. This can be configured for the main forwarding classes and for sub-classes, allowing each to use a different forwarding class at the egress.

The buffering, queuing, policing and remarking operation at the ingress and egress remain unchanged. Egress reclassification is possible. The profile processing (in/out) is completely unaffected by overriding the forwarding class.

When used in conjunction with QPPB (QoS Policy Propagation Using BGP), a QPPB assigned forwarding class takes precedence over both the normal ingress forwarding class classification rules and any egress forwarding class overrides.

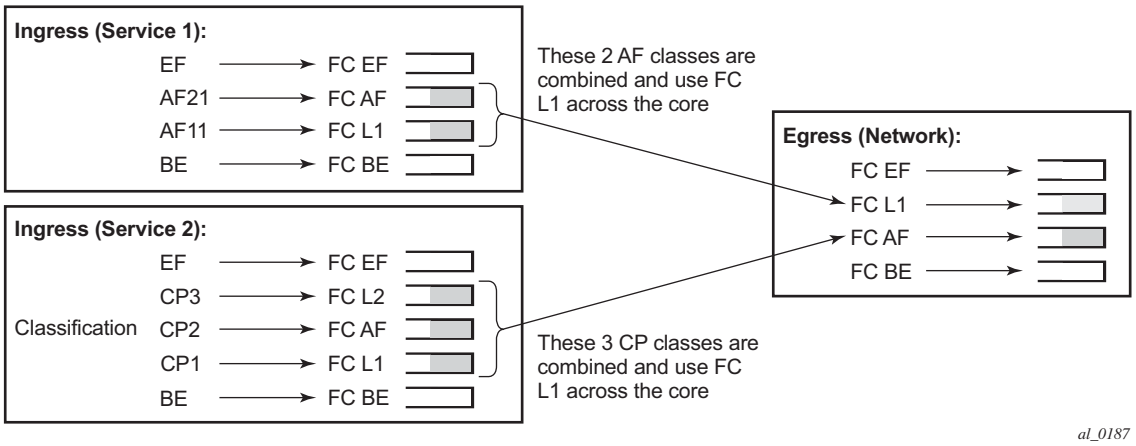

Figure 4: Egress Forwarding CLass Override

Figure 4 shows the ingress service 1 using forwarding classes AF and L1 that are overridden to L1 for the network egress, while it also shows ingress service 2 using forwarding classes L1, AF, and L2 that are overridden to AF for the network egress.

Service Egress QoS Policies

Service egress queues are implemented at the transition from the service core network to the service access network. The advantages of per-service queuing before transmission into the access network are:

- Per-service egress subrate capabilities especially for multipoint services.
- More granular, fairer scheduling per-service into the access network.
- Per-service statistics for forwarded and discarded service packets.

The subrate capabilities and per-service scheduling control are required to make multiple services per physical port possible. Without egress shaping, it is impossible to support more than one service per port. There is no way to prevent service traffic from bursting to the available port bandwidth and starving other services.

For accounting purposes, per-service statistics can be logged. When statistics from service ingress queues are compared with service egress queues, the ability to conform to per-service QoS requirements within the service core can be measured. The service core statistics are a major asset to core provisioning tools.

Service egress QoS policies define egress queues and map forwarding class flows to queues. In the simplest service egress QoS policy, all forwarding classes are treated like a single flow and mapped to a single queue. To define a basic egress QoS policy, the following are required:

- A unique service egress QoS policy ID.
- A QoS policy scope of template or exclusive.
- At least one defined default queue.

Optional service egress QoS policy elements include:

- Additional queues up to a total of 8 separate queues (unicast).
- IEEE 802.1p priority value remarking based on forwarding class.

Each queue in a policy is associated with one of the forwarding classes. Each queue can have its individual queue parameters allowing individual rate shaping of the forwarding class(es) mapped to the queue.

More complex service queuing models are supported in the router where each forwarding class is associated with a dedicated queue.

The forwarding class determination per service egress packet is determined at ingress. If the packet ingressed the service on the same router, the service ingress classification rules determine the forwarding class of the packet. If the packet is received, the forwarding class is marked in the tunnel transport encapsulation.

Service egress QoS policy ID 1 is reserved as the default service egress policy. The default policy cannot be deleted or changed. The default access egress policy is applied to all SAPs service egress policy explicitly assigned. The characteristics of the default policy are listed in the following table.

Slope Policies

For network ingress, a buffer pool is created for the and is used for all network ingress queues for ports on the .

Slope policies define the RED slope characteristics as a percentage of pool size for the pool on which the policy is applied.

Default buffer pools exist (logically) at the port and levels. Each physical port has two pools objects associated:

- Access ingress pool
- Access egress pool
- Network egress pool

By default, each pool is associated with slope-policy **default**.

Access, and network pools (in network mode) and access uplink pools (in access uplink mode) are created at the port level; creation is dependent on the physical port mode (network, access) or the mode of provisioned channel paths.

Node-level pools are used by ingress network queues and bundle access queues. A single ingress network pool is created at the node-level for ingress network queues.

An ingress and egress access pool is created at the MDA level for all bundle access queues.

RED Slopes

Operation and Configuration

Each buffer pool supports a high-priority RED slope, a non-TCP RED slope, and a low-priority RED slope. The high-priority RED slope manages access to the shared portion of the buffer pool for high-priority or in-profile packets. The low-priority RED slope manages access to the shared portion of the buffer pool for low-priority or out-of-profile packets.

For access buffer pools, the percentage of the buffers that are to be reserved for CBS buffers is configured by the user software (cannot be changed by user). This setting indirectly assigns the amount of shared buffers on the pool. This is an important function that controls the ultimate average and total shared buffer utilization value calculation used for RED slope operation. The CBS setting can be used to dynamically maintain the buffer space on which the RED slopes operate.

For network buffer pools, the CBS setting does not exist; instead, the configured CBS values for each network forwarding class queue inversely defines the shared buffer size. If the total CBS for each queue equals or exceeds 100% of the buffer pool size, the shared buffer size is equal to 0 (zero) and a queue cannot exceed its CBS.

When a queue depth exceeds the queue's CBS, packets received on that queue must contend with other queues exceeding their CBS for shared buffers. To resolve this contention, the buffer pool uses two RED slopes to determine buffer availability on a packet by packet basis. A packet that was either classified as high priority or considered in-profile is handled by the high-priority RED slope. This slope should be configured with RED parameters that prioritize buffer availability over packets associated with the low-priority RED slope. Packets that had been classified as low priority or out-of-profile are handled by this low-priority RED slope.

The following is a simplified overview of how a RED slope determines shared buffer availability on a packet basis:

1. The RED function keeps track of shared buffer utilization and shared buffer average utilization.
2. At initialization, the utilization is 0 (zero) and the average utilization is 0 (zero).
3. When each packet is received, the current average utilization is plotted on the slope to determine the packet's discard probability.
4. A random number is generated associated with the packet and is compared to the discard probability.
5. The lower the discard probability, the lower the chances are that the random number is within the discard range.

6. If the random number is within the range, the packet is discarded which results in no change to the utilization or average utilization of the shared buffers.
7. A packet is discarded if the utilization variable is equal to the shared buffer size or if the utilized CBS (actually in use by queues, not just defined by the CBS) is oversubscribed and has stolen buffers from the shared size, lowering the effective shared buffer size equal to the shared buffer utilization size.
8. If the packet is queued, a new shared buffer average utilization is calculated using the time-average-factor (TAF) for the buffer pool. The TAF describes the weighting between the previous shared buffer average utilization result and the new shared buffer utilization in determining the new shared buffer average utilization. (See Tuning the Shared Buffer Utilization Calculation on page 58.)
9. The new shared buffer average utilization is used as the shared buffer average utilization next time a packet's probability is plotted on the RED slope.
10. When a packet is removed from a queue (if the buffers returned to the buffer pool are from the shared buffers), the shared buffer utilization is reduced by the amount of buffers returned. If the buffers are from the CBS portion of the queue, the returned buffers do not result in a change in the shared buffer utilization.

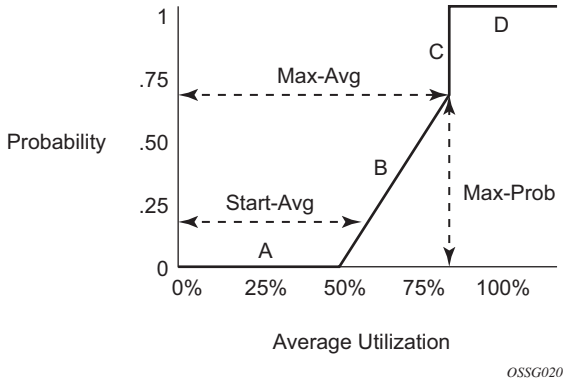

Figure 5: RED Slope Characteristics

A RED slope itself is a graph with an X (horizontal) and Y (vertical) axis. The X-axis plots the percentage of shared buffer average utilization, going from 0 to 100 percent. The Y-axis plots the probability of packet discard marked as 0 to 1. The actual slope can be defined as four sections in (X, Y) points (Figure 5):

1. Section A is (0, 0) to (start-avg, 0). This is the part of the slope that the packet discard value is always zero, preventing the RED function from discarding packets when the shared buffer average utilization falls between 0 and start-avg.
2. Section B is (start-avg, 0) to (max-avg, max-prob). This part of the slope describes a linear slope where packet discard probability increases from zero to max-prob.

3. Section C is (max-avg, max-prob) to (max-avg, 1). This part of the slope describes the instantaneous increase of packet discard probability from max-prob to one. A packet discard probability of 1 results in an automatic discard of the packet.
4. Section D is (max-avg, 1) to (100%, 1). On this part of the slope, the shared buffer average utilization value of max-avg to 100% results in a packet discard probability of 1.

Plotting any value of shared buffer average utilization will result in a value for packet discard probability from 0 to 1. Changing the values for start-avg, max-avg and max-prob allows the adaptation of the RED slope to the needs of the access or network queues using the shared portion of the buffer pool, including disabling the RED slope.

Tuning the Shared Buffer Utilization Calculation

The router allows tuning the calculation of the Shared Buffer Average Utilization (SBAU) after assigning buffers for a packet entering a queue as used by the RED slopes to calculate a packet's drop probability. The router implements a time average factor (TAF) parameter in the buffer policy which determines the contribution of the historical shared buffer utilization and the instantaneous Shared Buffer Utilization (SBU) in calculating the SBAU. The TAF defines a weighting exponent used to determine the portion of the shared buffer instantaneous utilization and the previous shared buffer average utilization used to calculate the new shared buffer average utilization. To derive the new shared buffer average utilization, the buffer pool takes a portion of the previous shared buffer average and adds it to the inverse portion of the instantaneous shared buffer utilization (SBU). The formula used to calculated the average shared buffer utilization is:

$$SBAU_n = \left(SBU \times \frac{1}{2^{TAF}}\right) + \left(SBAU_{n-1} \times \frac{2^{TAF} - 1}{2^{TAF}}\right)$$

where:

$SBAU_n$ = Shared buffer average utilization for event n

$SBAU_{n-1}$ = Shared buffer average utilization for event (n-1)

SBU = The instantaneous shared buffer utilization

TAF = The time average factor

Table 16 shows the effect the allowed values of TAF have on the relative weighting of the instantaneous SBU and the previous SBAU ($SBAU_{n-1}$) has on the calculating the current SBAU ($SBAU_n$).

Table 16: TAF Impact on Shared Buffer Average Utilization Calculation

TAF	2^{TAF}	Equates To	Shared Buffer Instantaneous Utilization Portion	Shared Buffer Average Utilization Portion
0	2^0	1	1/1 (1)	0 (0)
1	2^1	2	1/2 (0.5)	1/2 (0.5)
2	2^2	4	1/4 (0.25)	3/4 (0.75)
3	2^3	8	1/8 (0.125)	7/8 (0.875)
4	2^4	16	1/16 (0.0625)	15/16 (0.9375)

Table 16: TAF Impact on Shared Buffer Average Utilization Calculation (Continued)

TAF	2^{TAF}	Equates To	Shared Buffer Instantaneous Utilization Portion	Shared Buffer Average Utilization Portion
5	2^5	32	1/32 (0.03125)	31/32 (0.96875)
6	2^6	64	1/64 (0.015625)	63/64 (0.984375)
7	2^7	128	1/128 (0.0078125)	127/128 (0.9921875)
8	2^8	256	1/256 (0.00390625)	255/256 (0.99609375)
9	2^9	512	1/512 (0.001953125)	511/512 (0.998046875)
10	2^{10}	1024	1/1024 (0.0009765625)	1023/2024 (0.9990234375)
11	2^{11}	2048	1/2048 (0.00048828125)	2047/2048 (0.99951171875)
12	2^{12}	4096	1/4096 (0.000244140625)	4095/4096 (0.999755859375)
13	2^{13}	8192	1/8192 (0.0001220703125)	8191/8192 (0.9998779296875)
14	2^{14}	16384	1/16384 (0.00006103515625)	16383/16384 (0.99993896484375)
15	2^{15}	32768	1/32768 (0.000030517578125)	32767/32768 (0.999969482421875)

The value specified for the TAF affects the speed at which the shared buffer average utilization tracks the instantaneous shared buffer utilization. A low value weights the new shared buffer average utilization calculation more to the shared buffer instantaneous utilization. When TAF is zero, the shared buffer average utilization is equal to the instantaneous shared buffer utilization. A high value weights the new shared buffer average utilization calculation more to the previous shared buffer average utilization value. The TAF value applies to all high and low priority RED slopes for ingress and egress buffer pools controlled by the buffer policy.

Policy Parameters

The elements required to define a slope policy are:

- A unique policy ID
- The high and low RED slope shapes for the buffer pool: the start-avg, max-avg and max-prob.
- The TAF weighting factor to use for the SBAU calculation for determining RED slope drop probability.

Unlike access QoS policies where there are distinct policies for ingress and egress, slope policy is defined with generic parameters so that it is not inherently an ingress or an egress policy. A slope policy defines ingress properties when it is associated with an access port buffer pool on ingress and egress properties when it is associated with an access buffer pool on egress.

Each access port buffer pool can be associated with one slope policy ID on ingress and one slope policy ID on egress. The slope policy IDs on ingress and egress can be set independently.

Slope policy ID **default** is reserved for the default slope policy. The default policy cannot be deleted or changed. The default slope policy is implicitly applied to all access buffer pools which do not have another slope policy explicitly assigned.

Table 17 lists the default values for the default slope policy.

Table 17: Default Slope Policy Definition

Parameter	Description	Setting
Policy ID	Slope policy ID	1 (Policy ID 1 reserved for default slope policy)
High (RED) slope	Administrative state	Shutdown
	start-avg	70% utilization
	max-avg	90% utilization
	max-prob	80% probability
Low (RED) slope	Administrative state	Shutdown
	start-avg	50% utilization
	max-avg	75% utilization
	max-prob	80% probability
TAF	Time average factor	7

Table 18: Default Slope Policy Definition

Parameter	Description	Setting
Policy ID	Slope policy ID	1 (Policy ID 1 reserved for default slope policy)
High (RED) slope	Administrative state	Shutdown
	start-avg	70% utilization
	max-avg	90% utilization
	max-prob	80% probability
Low (RED) slope	Administrative state	Shutdown
	start-avg	50% utilization
	max-avg	75% utilization
	max-prob	80% probability
TAF	Time average factor	7

Scheduler Policies

A scheduler policy defines the hierarchy and all operating parameters for the member schedulers. A scheduler policy must be defined in the QoS context before a group of virtual schedulers can be used. Although configured in a scheduler policy, the individual schedulers are actually created when the policy is applied to a site, such as a SAP or interface.

Scheduler objects define bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. The scheduler object can also define a child association with a parent scheduler of its own.

A scheduler is used to define a bandwidth aggregation point within the hierarchy of virtual schedulers. The scheduler's rate defines the maximum bandwidth that the scheduler can consume. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can also be a child (take bandwidth from) a scheduler in a higher tier, except for schedulers created in Tier 1.

A parent parameter can be defined to specify a scheduler further up in the scheduler policy hierarchy. Only schedulers in Tiers 2 and 3 can have parental association. Tier 1 schedulers cannot have a parental association. When multiple schedulers and/or queues share a child status with the scheduler on the parent, the weight or strict parameters define how this scheduler contends with the other children for the parent's bandwidth. The parent scheduler can be removed or changed at anytime and is immediately reflected on the schedulers actually created by association of this scheduler policy.

When a parent scheduler is defined without specifying level, weight, or CIR parameters, the default bandwidth access method is weight with a value of 1.

If any orphaned queues (queues specifying a scheduler name that does not exist) exist on the ingress SAP and the policy application creates the required scheduler, the status on the queue becomes non-orphaned at this time.

Figure 6 depicts how child queues and schedulers interact with their parent scheduler to receive bandwidth. The scheduler distributes bandwidth to the children by first using each child's CIR according to the CIR-level parameter (CIR L8 through CIR L1 weighted loops). The weighting at each CIR-Level loop is defined by the CIR weight parameter for each child. The scheduler then distributes any remaining bandwidth to the children up to each child's rate parameter according to the Level parameter (L8 through L1 weighted loops). The weighting at each level loop is defined by the weight parameter for each child.

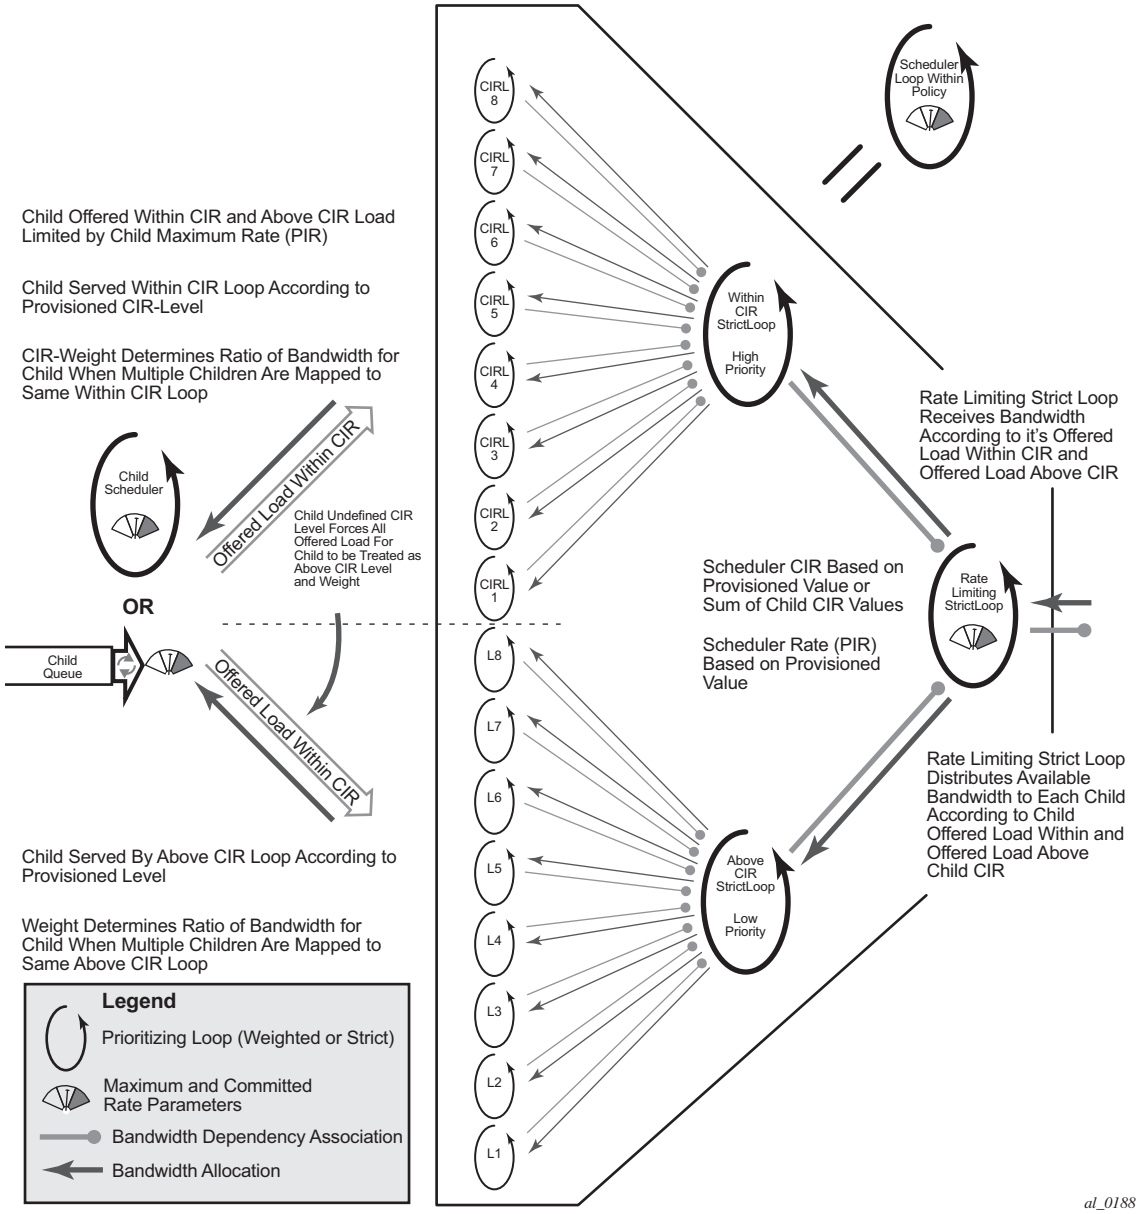

Figure 6: Virtual Scheduler Internal Bandwidth Allocation

Virtual Hierarchical Scheduling

Virtual hierarchical scheduling is a method that defines a bounded operation for a group of queues. One or more queues are mapped to a given scheduler with strict and weighted metrics controlling access to the scheduler. The scheduler has an optional prescribed maximum operating rate that limits the aggregate rate of the child queues. This scheduler may then feed into another virtual scheduler in a higher tier. The creation of a hierarchy of schedulers and the association of queues to the hierarchy allows for a hierarchical Service Level Agreement (SLA) to be enforced.

Scheduler policies in the routers determine the order queues are serviced. All ingress and egress queues operate within the context of a scheduler. Multiple queues share the same scheduler. Schedulers control the data transfer between the following queues and destinations:

- Service ingress queues to switch fabric destinations.
- Service egress queues to access egress ports.
- Network ingress queues to switch fabric destinations.
- Network egress queues to network egress interfaces.

There are two types of scheduler policies:

- Single Tier Scheduling on page 65
- Hierarchical Scheduler Policies on page 67

Schedulers and scheduler policies control the data transfer between queues, switch fabric destinations and egress ports/interfaces. The type of scheduling available for the various scheduling points within the system are summarized in Table 19.

Table 19: Supported Scheduler Policies

Scheduling From	To	Single-Tier	Hierarchical
Service Egress Queues	Access Egress Ports	Yes	Yes
Network Egress Queues	Network Egress Interfaces	Yes	No

Tiers

In single tier scheduling, queues are scheduled based on the forwarding class of the queue and the operational state of the queue relative to the queue's CIR and PIR. Queues operating within their CIR values are serviced before queues operating above their CIR values with "high-priority" forwarding class queues given preference over "low-priority" forwarding class queues. In single tier scheduling, all queues are treated as if they are at the same "level" and the queue's parameters and operational state directly dictate the queue's scheduling. Single tier scheduling is the system default scheduling policy for all the queues and destinations listed above and has no configurable parameters.

Hierarchical scheduler policies are an alternate way to schedule queues that can be used on service ingress and service egress queues. Hierarchical scheduler policies allow the creation of a hierarchy of schedulers where queues and/or other schedulers are scheduled by superior schedulers.

To illustrate the difference between single tier scheduling and hierarchical scheduling policies, consider a simple case where, on service ingress, three queues are created for gold, silver and bronze service and are configured as follows:

- Gold: CIR = 10 Mbps, PIR = 10 Mbps
- Silver: CIR = 20 Mbps, PIR = 40 Mbps
- Bronze: CIR = 0 Mbps, PIR = 100 Mbps

In the router, the CIR is used for policing of traffic (in-profile or out-of-profile), and the PIR is the rate at which traffic is shaped out of the queue. In single tier scheduling, each queue can burst up to its defined PIR, which means up to 150 Mbps (10 Mbps + 40 Mbps + 100 Mbps) can enter the service.

In a simple example of a hierarchical scheduling policy, a superior (or parent) scheduler can be created for the gold, silver and bronze queues which limits the overall rate for all queues to 100 Mbps. In this hierarchical scheduling policy, the customer can send in any combination of gold, silver and bronze traffic conforming to the defined PIR values and not to exceed 100 Mbps.

Single Tier Scheduling

Single-tier scheduling is the default method of scheduling queues in the router. Queues are scheduled with single-tier scheduling if no explicit hierarchical scheduler policy is defined or applied. There are no configurable parameters for single-tier scheduling.

In single tier scheduling, queues are scheduled based on the Forwarding Class of the queue and the operational state of the queue relative to the queue's Committed Information Rate (CIR) and Peak Information Rate (PIR). Queue's operating within their CIR values are serviced before queue's operating above their CIR values with "high-priority" forwarding class queues given preference over "low-priority" forwarding class queues. In Single Tier Scheduling, all queues are treated as if

they are at the same “level” and the queue’s parameters and operational state directly dictate the queue’s scheduling.

A pair of schedulers, a high-priority and low-priority scheduler, transmits to a single destination switch fabric port, access port, or network interface. Table 20 below lists how the forwarding class queues are mapped to the high and low scheduler:

Table 20: Forwarding Class Scheduler Mapping

Scheduler	Forwarding Class
High	Network Control
	Expedited
	High-2
	High 1
Low	Low-1
	Assured
	Low-2
	Best-Effort

Note, that by using the default QoS profile, all ingress traffic is treated as best effort (`be`) (mapped to FC `be` and to low priority scheduler). For an egress SAP using the default QoS profile, all egress traffic will use the same queue.

While competing for bandwidth to the destination, each scheduler determines which queue will be serviced next. During congestion (packets existing on multiple queues), queues are serviced in the following order:

1. Queues associated with the high-priority scheduler operating within their CIR.
2. Queues associated with the low-priority scheduler operating within their CIR.
3. All queues with traffic above CIR and within PIR will be serviced by a biased round robin.

Queues associated with a single scheduler are serviced in a round robin method. If a queue reaches the configured PIR, the scheduler will not serve the queue until the transmission rate drops below the PIR.

The router QoS features are flexible and allow modifications to the forwarding class characteristics and the CIR and PIR queue parameters. The only fundamental QoS mechanisms enforced within the hardware are the association of the forwarding classes with the high priority or low priority scheduler and the scheduling algorithm. Other parameters can be modified to configure the appropriate QoS behavior.

Hierarchical Scheduler Policies

Hierarchical scheduler policies are an alternate way of scheduling queues which can be used on service ingress and service egress queues. Hierarchical scheduler policies allow the creation of a hierarchy of schedulers where queues and/or other schedulers are scheduled by superior schedulers.

The use of the hierarchical scheduler policies is often referred to as hierarchical QoS or H-QoS on the SR OS.

Hierarchical Virtual Schedulers

Virtual schedulers are created within the context of a hierarchical scheduler policy. A hierarchical scheduler policy defines the hierarchy and parameters for each scheduler. A scheduler is defined in the context of a tier (Tier 1, Tier 2, Tier 3). The tier level determines the scheduler's position within the hierarchy. Three tiers of virtual schedulers are supported (Figure 7). Tier 1 schedulers (also called root schedulers) are defined without a parent scheduler. It is not necessary for Tier 1 schedulers to obtain bandwidth from a higher tier scheduler. A scheduler can enforce a maximum rate of operation for all child queues and associated schedulers.

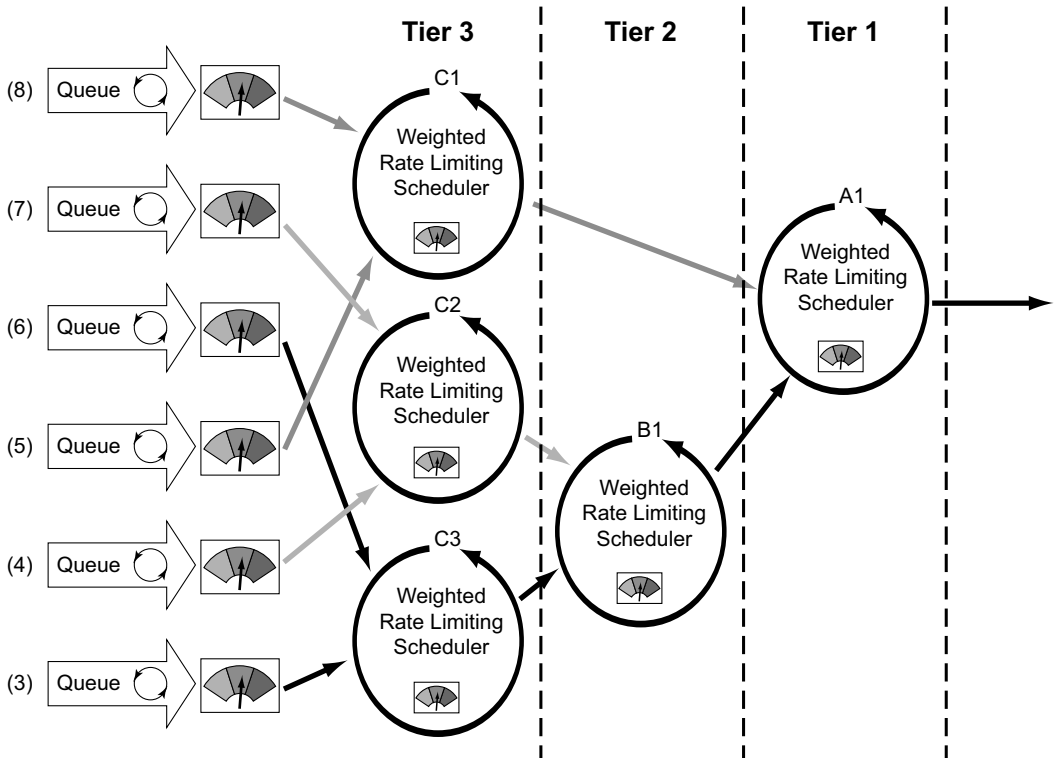

Figure 7: Hierarchical Scheduler and Queue Association

Scheduler Policies Applied to Applications

A scheduler policy can be applied either on a SAP (Figure 8) or on a multi-service customer site (a group of SAPs with common origination/termination point) (Figure 9). Whenever a scheduler policy is applied, the individual schedulers comprising the policy are created on the object. When the object is an individual SAP, only queues created on that SAP can use the schedulers created by the policy association. When the object is a multi-service customer site, the schedulers are available to any SAPs associated with the site (also see Scheduler Policies Applied to SAPs on page 70).

Refer to the Subscriber Services Overview section of the Services Guide for information about subscriber services, service entities, configuration, and implementation.

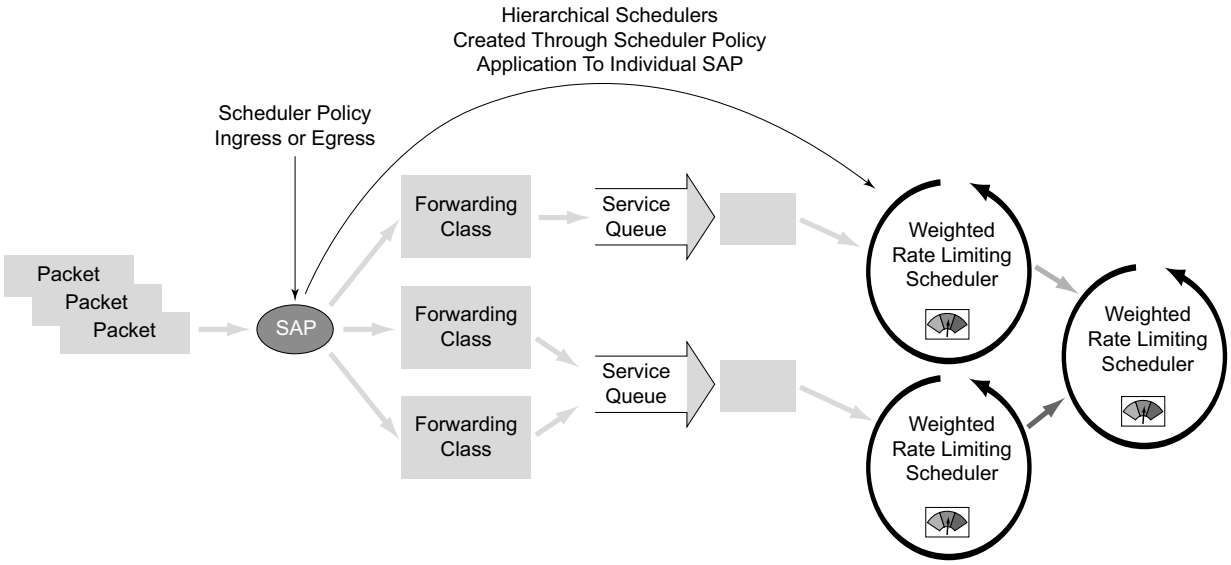

Figure 8: Scheduler Policy on SAP and Scheduler Hierarchy Creation

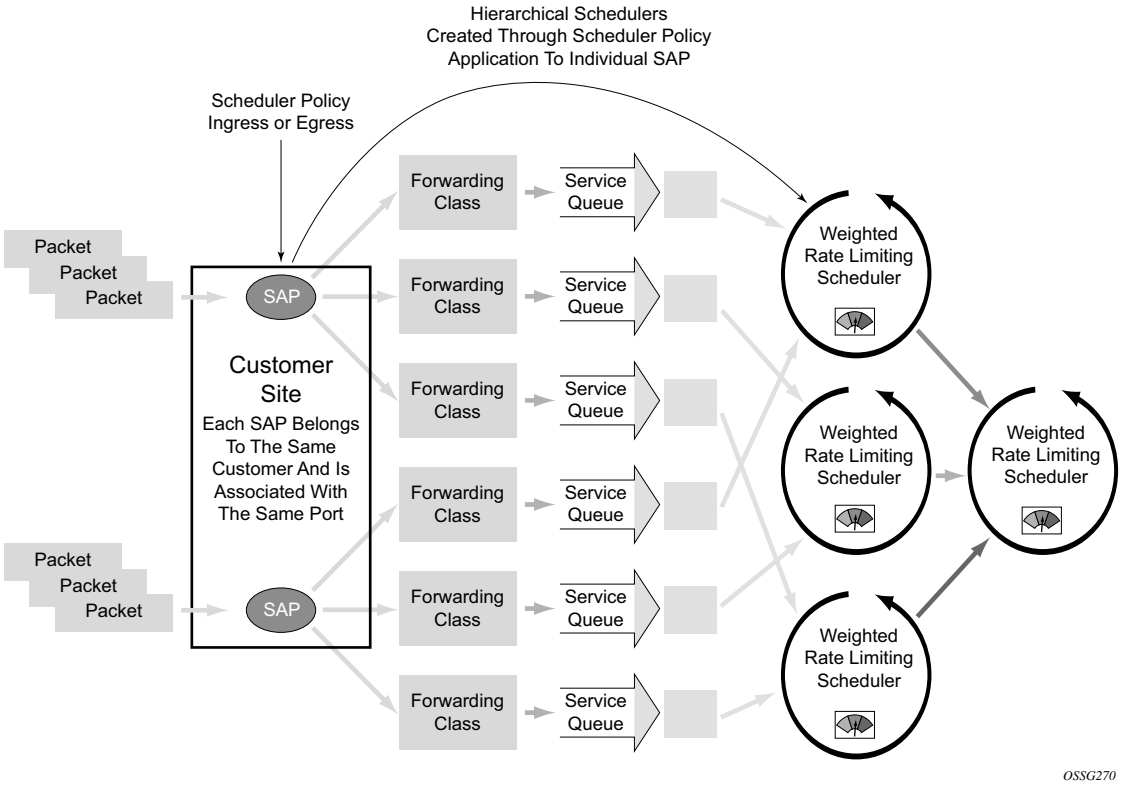

Figure 9: Scheduler Policy on Customer Site and Scheduler Hierarchy Creation

Queues become associated with schedulers when the parent scheduler name is defined within the queue definition in the SAP ingress policy. The scheduler is used to provide bandwidth to the queue relative to the operating constraints imposed by the scheduler hierarchy.

Scheduler Policies Applied to SAPs

A scheduler policy can be applied to create egress schedulers used by SAP queues. The schedulers comprising the policy are created at the time the scheduler policy is applied to the SAP. If any orphaned queues exist (queues specifying a scheduler name that does not exist) on the egress SAP and the policy application creates the required scheduler, the status on the queue will become non-orphaned.

Queues are associated with the configured schedulers by specifying the parent scheduler defined within the queue definition from the SAP egress policy. The scheduler is used to provide bandwidth to the queue relative to the operating constraints imposed by the scheduler hierarchy.

Customer Service Level Agreement (SLA)

The router implementation of hierarchical QoS allows a common set of virtual schedulers to govern bandwidth over a set of customer services that is considered to be from the same site. Different service types purchased from a single customer can be aggregately accounted and billed based on a single Service Level Agreement.

By configuring multi-service sites within a customer context, the customer site can be used as an anchor point to create an ingress and egress virtual scheduler hierarchy.

Once a site is created, it must be assigned to the chassis slot or a port . This allows the system to allocate the resources necessary to create the virtual schedulers defined in the ingress and egress scheduler policies. This also acts as verification that each SAP assigned to the site exists within the context of the customer ID and that the SAP was created on the correct slot, port, or channel. The specified slot or port must already be pre-provisioned (configured) on the system.

When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites are configured only to create a virtual scheduler hierarchy and make it available to queues on multiple SAPs.

Scheduler Policies Applied to Multi-Service Sites

Only an existing scheduler policy and scheduler policy names can be applied to create the ingress or egress schedulers used by SAP queues associated with a customer’s multi-service site. The schedulers defined in the scheduler policy can only be created after the customer site has been appropriately assigned to a chassis port, channel, or slot. Once a multi-service customer site is created, SAPs owned by the customer must be explicitly included in the site. The SAP must be owned by the customer the site was created within and the site assignment parameter must include the physical locale of the SAP.

Forwarding Classes

Routers support multiple forwarding classes and class-based queuing, so the concept of forwarding classes is common to all of the QoS policies.

Each forwarding class (also called Class of Service (CoS)) is important only in relation to the other forwarding classes. A forwarding class provides network elements a method to weigh the relative importance of one packet over another in a different forwarding class.

Queues are created for a specific forwarding class to determine the manner in which the queue output is scheduled into the switch fabric. The forwarding class of the packet, along with the in-profile or out-of-profile state, determines how the packet is queued and handled (the per hop behavior (PHB)) at each hop along its path to a destination egress point. Routers support eight (8) forwarding classes (Table 21).

Table 21: Forwarding Classes

FC-ID	FC Name	FC Designa-tion	DiffServ Name	Class Type	Notes
7	Network Control	NC	NC2	High-Priority	Intended for network control traffic.
6	High-1	H1	NC1		Intended for a second network control class or delay/jitter sensitive traffic.
5	Expedited	EF	EF		Intended for delay/jitter sensitive traffic.
4	High-2	H2	AF4		Intended for delay/jitter sensitive traffic.
3	Low-1	L1	AF2	Assured	Intended for assured traffic. Also is the default priority for network management traffic.
2	Assured	AF	AF1		Intended for assured traffic.
1	Low-2	L2	CS1	Best Effort	Intended for BE traffic.
0	Best Effort	BE	BE		

Note that Table 21 presents the default definitions for the forwarding classes. The forwarding class behavior, in terms of ingress marking interpretation and egress marking, can be changed by a Network QoS Policies on page 29. All forwarding class queues support the concept of in-profile and out-of-profile.

The forwarding classes can be classified into three class types:

- High-priority/Premium
- Assured
- Best effort

High-Priority Classes

The high-priority forwarding classes are Network Control (`nc`), Expedited (`ef`), High 1 (`h1`), and High 2 (`h2`). High-priority forwarding classes are always serviced at congestion points over other forwarding classes; this behavior is determined by the router queue scheduling algorithm (Virtual Hierarchical Scheduling on page 64).

With a strict PHB at each network hop, service latency is mainly affected by the amount of high-priority traffic at each hop. These classes are intended to be used for network control traffic or for delay or jitter-sensitive services.

If the service core network is over-subscribed, a mechanism to traffic engineer a path through the core network and reserve bandwidth must be used to apply strict control over the delay and bandwidth requirements of high-priority traffic. In the router, RSVP-TE can be used to create a path defined by an MPLS LSP through the core. Premium services are then mapped to the LSP with care exercised to not oversubscribe the reserved bandwidth.

If the core network has sufficient bandwidth, it is possible to effectively support the delay and jitter characteristics of high-priority traffic without utilizing traffic engineered paths, as long as the core treats high-priority traffic with the proper PHB.

Assured Classes

The assured forwarding classes are Assured (`af`) and Low 1 (`l1`). Assured forwarding classes provide services with a committed rate and a peak rate much like Frame Relay. Packets transmitted through the queue at or below the committed transmission rate are marked in-profile. If the core service network has sufficient bandwidth along the path for the assured traffic, all aggregate in-profile service packets will reach the service destination. Packets transmitted out the service queue that are above the committed rate will be marked out-of-profile. When an assured out-of-profile service packet is received at a congestion point in the network, it will be discarded before in-profile assured service packets.

Multiple assured classes are supported with relative weighting between them. In DiffServ, the code points for the various Assured classes are AF4, AF3, AF2 and AF1. Typically, AF4 has the highest weight of the four and AF1 the lowest. The Assured and Low 1 classes are differentiated based on the default DSCP mappings. Note that all DSCP and EXP mappings can be modified by the user.

Best-Effort Classes

The best-effort classes are Low 2 (`l2`) and Best-Effort (`be`). The best-effort forwarding classes have no delivery guarantees. All packets within this class are treated, at best, like out-of-profile assured service packets.

Shared Queues

Shared-queue QoS policies can be implemented to facilitate queue consumption on an MDA. It is especially useful when VPLS, IES, and VPRN services are scaled on one MDA. Instead of allocating multiple hardware queues for each unicast queue defined in a SAP ingress QoS policy, SAPs with the shared-queuing feature enabled only allocate one hardware queue for each SAP ingress QoS policy unicast queue.

However, as a tradeoff, the total amount of traffic throughput at ingress of the node is reduced because any ingress packet serviced by a shared-queuing SAP is recirculated for further processing. When the node is only used for access SAPs, 5 Gbps ingress traffic is the maximum that can be processed without seeing packet drops at the MDA ingress. The reason for this is that any ingress packet serviced by a shared-queuing SAP is processed twice in Flexible Fast Path which greatly reduces bandwidth.

Shared-queuing can add latency. Network planners should consider these restrictions while trying to scale services on one MDA.

ATM Traffic Descriptor Profiles

Traffic descriptors profiles capture the cell arrival pattern for resource allocation. Source traffic descriptors for an ATM connection include at least one of the following:

- Sustained Information Rate (SIR)
- Peak Information Rate (PIR)
- Minimum Information Rate (MIR)
- Maximum Burst Size (MBS)

QoS Traffic descriptor profiles are applied on IES, VPRN, VPLS, and VLL SAPs.

QoS Policy Entities

Services are configured with default QoS policies. Additional policies must be explicitly created and associated. There is one default service ingress QoS policy, one default service egress QoS policy, and one default network QoS policy. Only one ingress QoS policy and one egress QoS policy can be applied to a SAP or network port.

When you create a new QoS policy, default values are provided for most parameters with the exception of the policy ID and queue ID values, descriptions, and the default action queue assignment. Each policy has a scope, default action, a description, and at least one queue. The queue is associated with a forwarding class.

QoS policies can be applied to the following service types:

- Epipe — Both ingress and egress policies are supported on an Epipe service access point (SAP).
- VPLS — Both ingress and egress policies are supported on a VPLS SAP.
- IES — Both ingress and egress policies are supported on an IES SAP.
- VPRN — Both ingress and egress policies are supported on a VPRN SAP.

QoS policies can be applied to the following entities:

- Network ingress interface
- Network egress interface

Default QoS policies maps all traffic with equal priority and allow an equal chance of transmission (Best Effort (`be`) forwarding class) and an equal chance of being dropped during periods of congestion. QoS prioritizes traffic according to the forwarding class and uses congestion management to control access ingress, access egress, and network traffic with queuing according to priority

Frequently Used QoS Terms

The following terms are used in router Hierarchical QoS to describe the operation and maintenance of a virtual scheduler hierarchy and are presented for reference purposes.

Above CIR Distribution

'Above CIR' distribution is the second phase of bandwidth allocation between a parent scheduler and its child queues and child schedulers. The bandwidth that is available to the parent scheduler after the 'within CIR' distribution is distributed among the child members using each child's level (to define strict priority for the above CIR distribution), Weight (the ratio at a given level with several children) and the child's rate value. A rate value equal to the child's CIR value results in a child not receiving any bandwidth during the 'above CIR' distribution phase.

Available Bandwidth

Available bandwidth is the bandwidth usable by a parent scheduler to distribute to its child queues and schedulers. The available bandwidth is limited by the parent's schedulers association with its parent scheduler. If the parent scheduler has a parent of its own and the parent schedulers defined rate value, then available bandwidth is distributed to the child queues and schedulers using a 'within CIR' distribution phase and an 'above CIR' distribution phase. Distribution in each phase is based on a combination of the strict priority of each child and the relative weight of the child at that priority level. Separate priority and weight controls are supported per child for each phase.

CBS

The Committed Burst Size (CBS) specifies the relative amount of reserved buffers for a specific ingress network forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.

CIR

The Committed Information Rate (CIR) defines the amount of bandwidth committed to the scheduler or queue.

- For schedulers, the CIR value can be explicitly defined or derived from summing the child member CIR values.
- On a queue, the CIR value is explicitly defined.

  The CIR rate for ingress queues controls the in-profile and out-of-profile policing and ultimately egress in-profile and out-of-profile marking. Queue CIR rates also define the hardware fairness threshold at which the queue is no longer prioritized over other queues.

A child's (queue or scheduler) CIR is used with the CIR level parameter to determine the child's committed bandwidth from the parent scheduler. When multiple children are at the same strict CIR level, the CIR weight further determines the bandwidth distribution at that level.

CIR Level The CIR level parameter defines the strict level at which bandwidth is allocated to the child queue or scheduler during the within CIR distribution phase of bandwidth allocation. All committed bandwidth (determined by the CIR defined for the child) is allocated before any child receives non-committed bandwidth. Bandwidth is allocated to children at the higher CIR levels before children at a lower level. A child CIR value of zero or an undefined CIR level results in bandwidth allocation to the child only after all other children receive their provisioned CIR bandwidth. When multiple children share a CIR level, the CIR weight parameter further defines bandwidth allocation according to the child's weight ratio.

CIR Weight The CIR weight parameter defines the weight within the CIR level given to a child queue or scheduler. When multiple children share the same CIR level on a parent scheduler, the ratio of bandwidth given to an individual child is dependent on the ratio of the weights of the active children. A child is considered active when a portion of the offered load is within the child's defined CIR rate. The ratio is calculated by first adding the CIR weights of all active children and then dividing each child's CIR weight by the sum. If a child's CIR level parameter is not defined, that child is not included in the within CIR distribution and the CIR weight parameter is ignored. A CIR weight of zero forces the child to receive bandwidth only after all other children at that level have received their 'within CIR' bandwidth. When several children share a CIR weight of zero, all are treated equally.

Child Child is a logical state of a queue or scheduler that has been configured with a valid parent scheduler association. The child/parent association is used to build the hierarchy among the queues and schedulers.

Level The level parameter defines the strict priority level for a child queue or scheduler with regards to bandwidth allocation during the above CIR distribution phase on the child's parent scheduler. This allocation of bandwidth is done after the 'within CIR' distribution is finished. All child queues and schedulers receive the remaining bandwidth according to the strict priority level in which they are defined with higher levels receiving bandwidth first and lower levels receiving bandwidth if available.

MBS The Maximum Burst Size (MBS) command specifies the relative amount of the buffer pool space for the maximum buffers for a specific ingress network forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.

MCR The Minimum Cell Rate (MCR).

Offered Load Offered load is evaluated per child in the scheduler hierarchy. The offered load is the amount of bandwidth a child queue or scheduler can use to accommodate the data passing through the child. It is separated into two portions; within CIR and above CIR. Within CIR offered load is the portion of bandwidth required to meet the child's CIR value. It can be less than the CIR value but never

greater. If the forwarding requirement for the child is greater than the CIR value, the remaining is considered to be the above CIR offered load. The sum of the within CIR and above CIR offered load cannot be greater than the maximum rate defined for the child.

Orphan When a child queue is configured with a parent scheduler specified but the parent scheduler does not exist on the object the queue is created on, the state is considered orphaned.

An orphaned state is not the same condition as when a queue is not defined with a parent association. Orphan states are cleared when the parent scheduler becomes available on the object. This can occur when a scheduler policy containing the parent scheduler name is applied to the object that the queue exists on or when the scheduler name is added to the scheduler policy already applied to the object that the queue exists on.

Parent A scheduler becomes a parent when a queue or scheduler defines it as its parent. A queue or scheduler can be a child of only one scheduler. When defining a parent association on a child scheduler, the parent scheduler must already exist in the same scheduler policy and on a scheduler tier higher (numerically lower) then the child scheduler. Parent associations for queues are only checked once, when an instance of the queue is created on a SAP.

Queue A queue is where packets that will be forwarded are buffered before scheduling. Packets are not actually forwarded through the schedulers; they are forwarded from the queues directly to ingress or egress interfaces. The association between the queue and the virtual schedulers is intended to accomplish bandwidth allocation to the queue. Because the offered load is derived from queue utilization, bandwidth allocation is dependent on the queue distribution among the scheduler hierarchy. Queues can be tied to only one scheduler within the hierarchy.

Rate The rate defines the maximum bandwidth that will be made available to the scheduler or queue. The rate is defined in kilobits per second (Kbps).

- On a scheduler, the rate setting is used to limit the total bandwidth allocated to the scheduler's child members.
- For queues, the rate setting is used to define the Peak Information Rate (PIR) at which the queue can operate.

Root (Scheduler) A scheduler that has no parent scheduler association (is not a child of another scheduler) is considered to be a root scheduler. With no parent scheduler, bandwidth utilized by a root scheduler is dependent on offered load of child members, the maximum rate defined for the scheduler and total overall available bandwidth. Any scheduler can be a root scheduler. Since parent associations are not allowed in Tier 1, all schedulers in Tier 1 are considered be a root scheduler.

Scheduler Policy

A scheduler policy represents a particular grouping of virtual schedulers that are defined in specific scheduler tiers. The tiers and internal parent associations between the schedulers establish the hierarchy among the virtual schedulers. A scheduler policy can be applied to either a multi-service site or to a service Service Access Point (SAP). Once the policy is applied to a site or SAP, the schedulers in the policy are instantiated on the object and are available for use by child queues directly or indirectly associated with the object.

Tier

A tier is an organizational configuration used within a scheduler policy to define the place of schedulers created in the policy. Three tiers are supported; Tier 1, Tier 2, and Tier 3. Schedulers defined in Tier 2 can have parental associations with schedulers defined in Tier 1. Schedulers defined in Tier 3 can have parental associations with schedulers defined at Tiers 1 or 2. Queues can have parental associations with schedulers at any tier level.

Virtual Scheduler

A virtual scheduler, defined by a name (text string), is a logical configuration used as a parent to a group of child members that are dependent upon a common parent for bandwidth allocation. The virtual scheduler can also be a child member to another parent virtual scheduler and receive bandwidth from that parent to distribute to its child members.

Weight

The weight parameter defines the weight within the ‘above CIR’ level given to a child queue or scheduler. When several children share the same level on a parent scheduler, the ratio of bandwidth give to an individual child is dependent on the ratio of the weights of the active children. A child is considered active when a portion of the offered load is above the CIR value (also bounded by the child’s maximum bandwidth defined by the child’s rate parameter). The portion of bandwidth given to each child is based on the child’s weight compared to the sum of the weights of all active children at that level. A weight of zero forces the child to receive bandwidth only after all other children at that level have received their ‘above CIR’ bandwidth. When several children share a weight of zero, all are treated equally.

Within CIR Distribution

Within the CIR distribution process is the initial phase of bandwidth allocation between a parent scheduler and its child queues and child schedulers. The bandwidth that is available to the parent scheduler is distributed first among the child members using each child’s CIR level (to define a strict priority for the CIR distribution), CIR weight (the ratio at a given CIR level with several children), and the child’s CIR value. A CIR value of zero or an undefined CIR level causes a child to not receive any bandwidth during the CIR distribution phase. If the parent scheduler has any bandwidth remaining after the ‘within CIR’ distribution phase, it will be distributed using the above CIR distribution phase.

Configuration Notes

The following information describes QoS implementation caveats:

- Creating additional QoS policies is optional.
- Default policies are created for service ingress, service egress, network, network-queue, slope policies. Scheduler policies must be explicitly created and applied to a port.
- Associating a service or access ports with a QoS policy other than the default policy is optional.
- A network queue, service egress, and service ingress QoS policy must consist of at least one queue. Queues define the forwarding class, CIR, and PIR associated with the queue.

Network QoS Policies

In This Section

This section provides information to configure network QoS policies using the command line interface.

Topics in this section include:

- Overview on page 82
- Basic Configurations on page 89
- Default Network Policy Values on page 92
- Service Management Tasks on page 97

Overview

The ingress component of the policy defines how DiffServ code points (DSCPs) and MPLS EXP bits are mapped to internal forwarding class and profile state. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the router. The mapping on each network interface defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the network interface.

The egress component of the network QoS policy defines the DiffServ oriented queuing parameters associated with each forwarding class.

Each forwarding class defined within the system automatically creates a queue on each network interface. This queue gets all the parameters defined within the default network QoS policy 1 until an explicit policy is defined for the network interface.

If the egressing packet originated on an ingress SAP, or the remarking parameter is defined for the egress interface, the egress QoS policy also defines the IP DSCP or MPLS EXP bit marking based on the forwarding class and the profile state.

Network **policy-id 1** exists as the default policy that is applied to all network interfaces by default. The network **policy-id 1** cannot be modified or deleted. It defines the default DSCP-to-FC mapping and MPLS EXP-to-FC for the ingress. For the egress, it defines six forwarding classes which represent individual queues and the packet marking criteria.

New (non-default) network policy parameters can be modified. The **no** form of the command reverts the object to the default values. A new network policy must include the definition of at least one queue and specify the default-action. Incomplete network policies cannot be applied to network interfaces.

Changes made to a policy are applied immediately to all network interface where the policy is applied. For this reason, when a policy requires several changes, it is recommended that you copy the policy to a work area policy-id. The work-in-progress copy can be modified until all the changes are made and then the original policy-id can be overwritten with the **config qos copy** command.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your router devices, refer to CLI Usage chapter in the Basic System Configuration Guide.

Network Ingress Tunnel QoS Override

For Tunnel Terminated IP Routing Decisions

This section describes a mechanism that provides the ability to ignore the network ingress QoS mapping of a terminated tunnel containing an IP packet that is to be routed to a base router or VPRNdestination. This is advantageous when the mapping for the tunnel QoS marking does not accurately or completely reflect the required QoS handling for the IP routed packet. When the mechanism is enabled on an ingress network IP interface, the IP interface will ignore the tunnel's QoS mapping and derive the internal forwarding class and profile based on the precedence or DiffServe Code Point (DSCP) values within the routed IP header ToS field compared to the Network QoS policy defined on the IP interface.

Normal QoS Operation

The following types of QoS mapping decisions are applicable on a network ingress IP interface.

- Ethernet dot1p value mapping (if defined)
- Default QoS mapping
- IP ToS precedence mapping
- IP ToS DSCP mapping
- MPLS LSP EXP mapping

The default QoS mapping always exists on an ingress IP interface and every received packet will be mapped to this default if another explicitly defined matching entry does not exist.

A tunnel that terminates on the ingress IP interface (the node is the last hop for the tunnel) is evaluated based on the type of tunnel, IP GRE or MPLS LSP. An IP tunneled packet may match a dot1p entry, IP ToS precedence entry or IP ToS DSCP entry when defined in the applied policy. An MPLS LSP may match a dot1p entry or MPLS EXP entry when defined.

The internal tunnel encapsulated packet is never evaluated for QoS determination when operating in normal mode.

Tunnel Termination QoS Override Operation

Tunnel termination QoS override only applies to IP routing decisions once the tunnel encapsulation is removed. Non-IP routed packets within a terminating tunnel are ignored by the override and are forwarded as described in the Normal QoS Operation section.

When tunnel termination QoS override is enabled, the ToS field within the routed IP header is evaluated against the IP ToS precedence and DSCP entries in the applied network QoS policy on the ingress IP interface. If an explicit match entry is not found, the default QoS mapping is used. Any dot1p and MPLS LSP EXP bits within the packet are ignored. If the packet was IP GRE tunneled to the node, the tunnel IP header ToS field is ignored as well.

Any tunnel received on the ingress IP interface that traverses the node (the node is not the ultimate hop for the tunnel) is not affected by the QoS override mechanism and is forwarded as described in Normal QoS Operation section.

Enabling and Disabling Tunnel Termination QoS Override

Tunnel termination QoS override is enabled and disabled within the network QoS policy under the ingress node. The default condition within the policy is not to override tunnel QoS for IP routed packets.

DSCP Marking CPU Generated Traffic

Specific DSCP, forwarding class, and dot1p parameters can be specified to be used by every protocol packet generated by the node. This enables prioritization or de-prioritization of every protocol (as required). The markings effect a change in behavior on ingress when queuing. For example, if OSPF is not enabled, then traffic can be de-prioritized to best effort (be) DSCP. This change de-prioritizes OSPF traffic to the CPU complex.

DSCP marking for internally generated control and management traffic by marking the DSCP value should be used for the given application. This can be configured per routing instance. For example, OSPF packets can carry a different DSCP marking for the base instance and then for a VPRN service. IS-IS and ARP traffic is not an IP-generated traffic type and is not DSCP configurable, but they are dot1p configurable.

When an application is configured to use a specified DSCP value then the MPLS EXP, dot1p bits will be marked in accordance with the network or access egress policy as it applies to the logical interface the packet will be egressing.

The DSCP value can be set per application. This setting will be forwarded to the egress IOM. The egress IOM does not alter the coded DSCP value and marks the LSP-EXP and IEEE 802.1p (dot1p) bits according to the appropriate network or access QoS policy.

Sgt-qos is supported in the base router, VPRN and management contexts.

The default values for self-generated traffic are:

- Routing protocols (OSPF, BGP, etc)
  - → Forwarding class: Network Control (NC)
  - → DSCP value: NC1 (not applicable for ARP, IS-IS and PPPoE)
  - → 802.1p value: 7
- Management protocols (SSH, SNMP, etc)
  - → Forwarding class: Assured Forwarding (AF)
  - → DSCP value: AF41

Table 22: DSCP/FC Marking

Protocol	IPv4	IPv6	DSCP Marking	dot1p Marking	Default FC
IS-IS				Yes	NC
ARP				Yes	NC
PPPoE				Yes	NC

Table 22: DSCP/FC Marking (Continued)

Protocol	IPv4	IPv6	DSCP Marking	dot1p Marking	Default FC
LDP (T-LDP)	Yes				NC
RSVP	Yes		Yes	Yes	NC
BGP	Yes	Yes	Yes	Yes	NC
RIP	Yes	Yes	Yes	Yes	NC
MSDP	Yes				NC
PIM (SSM)	Yes	Yes	Yes	Yes	NC
OSPF	Yes	Yes	Yes	Yes	NC
SMTP	Yes				AF
IGMP/MLD	Yes	Yes	Yes	Yes	AF
Telnet	Yes	Yes	Yes	Yes	AF
TFTP	Yes		Yes	Yes	AF
FTP	Yes				AF
SSH (SCP)	Yes	Yes	Yes	Yes	AF
SNMP (get, set, etc.)	Yes	Yes	Yes	Yes	AF
SNMP trap/log	Yes	Yes	Yes	Yes	AF
syslog	Yes	Yes	Yes	Yes	AF
ICMP	Yes	Yes	Yes	Yes	AF
Traceroute	Yes	Yes	Yes	Yes	AF
TACPLUS	Yes	Yes	Yes	Yes	AF
DNS	Yes	Yes	Yes	Yes	AF
SNTP/NTP	Yes				AF
RADIUS	Yes				AF
Cflowd	Yes				AF
DHCP	Yes	Yes	Yes	Yes	AF
IPv6 Neighbor Discovery	Yes				NC

NOTE: The ICMP entry under sgt-qos is not referenced for ICMP ECHO_REQUEST (8) and ECHO_RESPONSE (0) packet types.

Default DSCP Mapping Table

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary	Label
Default	0	0x00	0b000000	be
nc1	48	0x30	0b110000	h1
nc2	56	0x38	0b111000	nc
ef	46	0x2e	0b101110	ef
af11	10	0x0a	0b001010	assured
af12	12	0x0c	0b001100	assured
af13	14	0x0e	0b001110	assured
af21	18	0x12	0b010010	l1
af22	20	0x14	0b010100	l1
af23	22	0x16	0b010110	l1
af31	26	0x1a	0b011010	l1
af32	28	0x1c	0b011100	l1
af33	30	0x1d	0b011110	l1
af41	34	0x22	0b100010	h2
af42	36	0x24	0b100100	h2
af43	38	0x26	0b100110	h2
default*	0			

*The default forwarding class mapping is used for all DSCP names/values for which there is no explicit forwarding class mapping.

Basic Configurations

A basic network QoS policy must conform to the following:

- Each network QoS policy must have a unique policy ID.
- Include the definition of at least one queue.
- Specify the default-action.

Create a Network QoS Policy

Configuring and applying QoS policies other than the default policy is optional. A default network policy of the appropriate type is applied to each router interface.

To create an network QoS policy when operating, define the following:

- A network policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- You can modify egress criteria to customize the forwarding class queues to be instantiated. Otherwise, the default values are applied.
  - → Remarking — When enabled, this command remarks ALL packets that egress on the specified network port. The remarking is based on the forwarding class to DSCP and LSP EXP bit mapping defined under the egress node of the network QoS policy.
  - → Forwarding class criteria — The forwarding class name represents an egress queue. Specify forwarding class criteria to define the egress characteristics of the queue and the marking criteria of packets flowing through it.
  - → DSCP — The DSCP value is used for all IP packets requiring marking that egress on this forwarding class queue that are *in* or *out* of profile.
  - → LSP EXP — The EXP value is used for all MPLS labeled packets requiring marking that egress on this forwarding class queue that are *in* or *out* of profile.
- Ingress criteria — Specifies the DSCPdot1p to forwarding class mapping for all IP packets and define the MPLS EXP bits to forwarding class mapping for all labeled packets.
  - → Default action — Defines the default action to be taken for packets that have an undefined DSCP or MPLS EXP bits set. The default-action specifies the forwarding class to which such packets are assigned.
  - → DSCP — Creates a mapping between the DSCP of the network ingress traffic and the forwarding class. Ingress traffic that matches the specified DSCP will be assigned to the corresponding forwarding class.

→ LSP EXP — Creates a mapping between the LSP EXP bits of the network ingress traffic and the forwarding class. Ingress traffic that matches the specified LSP EXP bits will be assigned to the corresponding forwarding class.

Use the following CLI syntax to create a network QoS policy:

CLI Syntax:

```
config>qos#
    network network-policy-id
        description description-string
        scope {exclusive|template}
        egress
            remarking
            fc {be|l2|af|l1|h2|ef|h1|nc}
                dot1p-in-profile dot1p-priority
                dot1p-out-profile dot1p-priority
                dscp-in-profile dscp-name
                dscp-out-profile dscp-name
                lsp-exp-in-profile mpls-exp-value
                lsp-exp-out-profile mpls-exp-value
            default-action fc {be|l2|af|l1|h2|ef|h1|nc} profile
                {in|out}
            dot1p dot1p-priority fc {fc-name} profile {in|out}
            dscp dscp-name fc {be|l2|af|l1|h2|ef|h1|nc} profile
                {in|out}
            ler-use-dscp
            lsp-exp lsp-exp-value fc fc-name profile {in|out}
```

```
A:ALA-10:A:ALA-12>config>qos# info
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
     network 600 create
             description "Network Egress Policy"
             ingress
                 default-action fc ef profile in
             exit
             egress
                 remarking
             exit
         exit
...
#------------------------------------------
A:ALA-12>config>qos#
```

Applying Network Policies

Use the following CLI syntax to apply network policies to the routeraccess uplink portsIP interfaces:

CLI Syntax:
```
config>router
    interface interface-name
        qos network-policy-id
```

The following output displays the configuration for router interface ALA-1-2 with network policy 600 applied to the interface.

```
A:ALA-7>config>router# info
#------------------------------------------
echo "IP Configuration"
#------------------------------------------
...
        interface "ALA-1-2"
            address 10.10.4.3/24
            qos 600
        exit
...
-------------------------------------------------
A:ALA-7>config>router#
```

Default Network Policy Values

The default network policy for IP interfaces is identified as policy-id **1**. Default policies cannot be modified or deleted. The following displays default network policy parameters:

Table 23: Network Policy Defaults

Field	Default	
description	Default network QoS policy.	
scope	template	
ingress		
default-action	fc be profile out	
dscp:		
be	fc be	profile out
ef	fc ef	profile in
cs1	fc l2	profile in
nc1	fc h1	profile in
nc2	fc nc	profile in
af11	fc af	profile in
af12	fc af	profile out
af13	fc af	profile out
af21	fc l1	profile in
af22	fc l1	profile out
af23	fc l1	profile out
af31	fc l1	profile in
af32	fc l1	profile out
af33	fc l1	profile out
af41	fc h2	profile in
af42	fc h2	profile out

Table 23: Network Policy Defaults (Continued)

Field	Default	
af43	fc h2	profile out
lsp-exp:		
0	fc be	profile out
1	fc l2	profile in
2	fc af	profile out
3	fc af	profile in
4	fc h2	profile in
5	fc ef	profile in
6	fc h1	profile in
7	fc nc	profile in
egress		
remarking	no	
fc af:		
dscp-in-profile	af11	
dscp-out-profile	af12	
lsp-exp-in-profile	3	
lsp-exp-out-profile	2	
fc be:		
dscp-in-profile	be	
dscp-out-profile	be	
lsp-exp-in-profile	0	
lsp-exp-out-profile	0	
fc ef:		
dscp-in-profile	ef	
dscp-out-profile	ef	

Table 23: Network Policy Defaults (Continued)

Field	Default
lsp-exp-in-profile	5
lsp-exp-out-profile	5
fc h1:	
dscp-in-profile	nc1
dscp-out-profile	nc1
lsp-exp-in-profile	6
lsp-exp-out-profile	6
fc h2:	
dscp-in-profile	af41
dscp-out-profile	af42
lsp-exp-in-profile	4
lsp-exp-out-profile	4
fc l1:	
dscp-in-profile	af21
dscp-out-profile	af22
lsp-exp-in-profile	3
lsp-exp-out-profile	2
fc l2:	
dscp-in-profile	cs1
dscp-out-profile	cs1
lsp-exp-in-profile	1
lsp-exp-out-profile	1
fc nc:	
dscp-in-profile	nc2
dscp-out-profile	nc2
lsp-exp-in-profile	7

Table 23: Network Policy Defaults (Continued)

Field	Default
lsp-exp-out-profile	7

The following output displays the default configuration:

```
A:ALA-49>config>qos>network# info detail
----------------------------------------------
            description "Default network QoS policy."
            scope template
            ingress
                default-action fc be profile out
                no ler-use-dscp
                dscp be fc be profile out
                dscp ef fc ef profile in
                dscp cs1 fc l2 profile in
                dscp nc1 fc h1 profile in
                dscp nc2 fc nc profile in
                dscp af11 fc af profile in
                dscp af12 fc af profile out
                dscp af13 fc af profile out
                dscp af21 fc l1 profile in
                dscp af22 fc l1 profile out
                dscp af23 fc l1 profile out
                dscp af31 fc l1 profile in
                dscp af32 fc l1 profile out
                dscp af33 fc l1 profile out
                dscp af41 fc h2 profile in
                dscp af42 fc h2 profile out
                dscp af43 fc h2 profile out
                lsp-exp 0 fc be profile out
                lsp-exp 1 fc l2 profile in
                lsp-exp 2 fc af profile out
                lsp-exp 3 fc af profile in
                lsp-exp 4 fc h2 profile in
                lsp-exp 5 fc ef profile in
                lsp-exp 6 fc h1 profile in
                lsp-exp 7 fc nc profile in
            exit
            egress
                no remarking
                fc af
                    dscp-in-profile af11
                    dscp-out-profile af12
                    lsp-exp-in-profile 3
                    lsp-exp-out-profile 2
                    dot1p-in-profile 2
                    dot1p-out-profile 2
                exit
                fc be
                    dscp-in-profile be
                    dscp-out-profile be
                    lsp-exp-in-profile 0
```

```
                        lsp-exp-out-profile 0
                        dot1p-in-profile 0
                        dot1p-out-profile 0
                    exit
                    fc ef
                        dscp-in-profile ef
                        dscp-out-profile ef
                        lsp-exp-in-profile 5
                        lsp-exp-out-profile 5
                        dot1p-in-profile 5
                        dot1p-out-profile 5
                    exit
                    fc h1
                        dscp-in-profile nc1
                        dscp-out-profile nc1
                        lsp-exp-in-profile 6
                        lsp-exp-out-profile 6
                        dot1p-in-profile 6
                        dot1p-out-profile 6
                    exit
                    fc h2
                        dscp-in-profile af41
                        dscp-out-profile af42
                        lsp-exp-in-profile 4
                        lsp-exp-out-profile 4
                        dot1p-in-profile 4
                        dot1p-out-profile 4
                    exit
                    fc l1
                        dscp-in-profile af21
                        dscp-out-profile af22
                        lsp-exp-in-profile 3
                        lsp-exp-out-profile 2
                        dot1p-in-profile 3
                        dot1p-out-profile 3
                    exit
                    fc l2
                        dscp-in-profile cs1
                        dscp-out-profile cs1
                        lsp-exp-in-profile 1
                        lsp-exp-out-profile 1
                        dot1p-in-profile 1
                        dot1p-out-profile 1
                    exit
                    fc nc
                        dscp-in-profile nc2
                        dscp-out-profile nc2
                        lsp-exp-in-profile 7
                        lsp-exp-out-profile 7
                        dot1p-in-profile 7
                        dot1p-out-profile 7
                    exit
                exit
------------------------------------------------
A:ALA-49>config>qos>network#
```

Service Management Tasks

Deleting QoS Policies

A network policy is associated by default with router interfaces.

You can replace the default policy with a non-default policy, but you cannot remove default policies from the configuration. When you remove a non-default policy, the policy association reverts to the appropriate default network policy.

CLI Syntax:
```
config>router
    interface interface-name
        qos network-policy-id
```

The following output displays a sample configuration.

```
A:ALA-7>config>router# info
#------------------------------------------
echo "IP Configuration"
#------------------------------------------
...
        interface "ALA-1-2"
            address 10.10.4.3/24 broadcast host-ones
            no port
            no arp-timeout
            no allow-directed-broadcasts
            icmp
                mask-reply
                redirects 100 10
                unreachables 100 10
                ttl-expired 100 10
            exit
            qos 1
            ingress
                no filter
            exit
            egress
                no filter
            exit
            no mac
            no ntp-broadcast
            no cflowd
            no shutdown
        exit
        interface "ALA-1-3"
...
#------------------------------------------
A:ALA-7>config>router#
```

Remove a Policy from the QoS Configuration

To delete a network policy, enter the following commands:

CLI Syntax: `config>qos# no network` *`network-policy-id`*

Copying and Overwriting Network Policies

You can copy an existing network policy to a new policy ID value or overwrite an existing policy ID. The `overwrite` option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos# copy network` *`source-policy-id dest-policy-id`* `[overwrite]`

The following output displays the copied policies:

```
A:ALA-12>config>qos# info detail
----------------------------------------------
...
        network 1 create
            description "Default network QoS policy."
            scope template
            ingress
                default-action fc be profile out
                dscp be fc be profile out
                dscp ef fc ef profile in
                dscp cs1 fc l2 profile in
                dscp nc1 fc h1 profile in
                dscp nc2 fc nc profile in
                dscp af11 fc af profile in
                dscp af12 fc af profile out
                dscp af13 fc af profile out
                dscp af21 fc l1 profile in
                dscp af22 fc l1 profile out
...
        network 600 create
            description "Default network QoS policy."
            scope template
            ingress
                default-action fc be profile out
                dscp be fc be profile out
                dscp ef fc ef profile in
                dscp cs1 fc l2 profile in
                dscp nc1 fc h1 profile in
                dscp nc2 fc nc profile in
                dscp af11 fc af profile in
                dscp af12 fc af profile out
                dscp af13 fc af profile out
                dscp af21 fc l1 profile in
```

```
                dscp af22 fc l1 profile out
...
        network 700 create
            description "Default network QoS policy."
            scope template
            ingress
                default-action fc be profile out
                dscp be fc be profile out
                dscp ef fc ef profile in
                dscp cs1 fc l2 profile in
                dscp nc1 fc h1 profile in
                dscp nc2 fc nc profile in
                dscp af11 fc af profile in
                dscp af12 fc af profile out
                dscp af13 fc af profile out
                dscp af21 fc l1 profile in
                dscp af22 fc l1 profile out
...
-------------------------------------------
A:ALA-12>config>qos#
```

Editing QoS Policies

You can change existing policies, except the default policies, and entries in the CLI. The changes are applied immediately to all interfaces where the policy is applied. To prevent configuration errors use the copy command to make a duplicate of the original policy to a work area, make the edits, and then overwrite the original policy.

Resource Allocation for Network QoS policy

This section describes the allocation of QoS resources for network QoS policy (for type=ipinterface).

When an IP interface is created, a default network QoS policy is applied. For the default policy, two meters and two classification entries in hardware are allocated.

The resources are allocated to a network policy, only when a port is configured for the IP interface.

For every FC in use, the system allocates two classification entries in hardware. If multiple match criteria entries map to the same FC, then each of these are allocated two classification entries in hardware. For example, if there are two match-criteria entries that map to FC ‘af’, then a total of four classification entries are allocated in hardware and if there are four match-criteria entries that map to FC ‘af’, then a total of 8 classification entries are allocated in hardware.

For every meter or policer in use, the system allocates one meter in hardware. A meter or policer is considered to be in use when it is associated with an FC in use.

The number of IP interfaces allowed is limited to number of resources available in hardware, subject to system limit (a maximum of 32 IP interfaces are allowed). The system reserves a total of 512 classification entries and 256 meters in hardware for use by network policy associated with an IP interface.

For computing the number of QoS resources used by an IP interface:

- Determine number of match-criteria entries used to identify the FC.
- Determine number of FCs to use.

Only the FCs used by the match-criteria classification entries are to be considered for the 'number of FCs'. Therefore are referred to as ‘FC in use’.

Use the following rules to compute the number of classification entries per FC in use:

If a FC is in use and is created without explicit meters, use default meter#1 for unicast traffic and default meter #9 for all other traffic types (that is, broadcast, multicast and unknown-unicast). This requires two classification entries in hardware.

If a FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter #9 for all other traffic types. This requires two classification entries in hardware.

If a FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic. This requires two classification entries in hardware.

Given the number of match criteria and the number of FCs used, use the equation given below to arrive at total number of classification entries per policy (for example TC):

$$TC = \Sigma_{i=nc,h1,ef,h2,l1,af,l2,be} \; 2 * E(i)$$

Where,

E(i) is the number of match- criteria entries that classify packets to FCi. For 7210 platforms, the maximum number of classification entries per policy can be 64 (including default).

2 is the number of classification entries that are required by FCi.

Note: In any case, only 2 classification entries are used per FC in a network policy, as only two traffic-types are supported.

Determine number of policers or meters to use (for example TP). A maximum of 12 meters per network policy is available.

Only those meters that are associated with FCs need to be considered for number of meters. Note, that only FCs in use are considered.

Network QoS Policy Command Reference

Command Hierarchies

- Configuration Commands on page 103
- Multi-Class Frame-Relay Profile Commands on page 103
- Operational Commands on page 104
- Show Commands on page 104

Configuration Commands

Multi-Class Frame-Relay Profile Commands

```
config
    — qos
        — [no] mc-fr-profile-ingress profile-id
            — description description-string
            — no description
            — class class-id
                — reassembly-timeout timeout-value
                — no reassembly-timeout
        — [no] mc-fr-profile-egress profile-id
            — description description-string
            — no description
            — class class-id
                — max-queue-size queue-size
                — no max-queue-size
                — mir mir
                — no mir
                — weight weight
                — no weight
config
    — qos
        — [no] network network-policy-id
            — description description-string
            — no description
            — scope {exclusive | template}
            — no scope
            — egress
                — [no] fc fc-name
                    — de-mark [force de-value
                    — no de-mark
                    — dot1p dot1p-priority
                    — no dot1p
                    — dot1p-in-profile dot1p-priority
                    — no dot1p-in-profile
                    — dot1p-out-profile dot1p-priority
                    — no dot1p-out-profile
```

```
                        — dscp-in-profile dscp-name
                        — no dscp-in-profile
                        — dscp-out-profile dscp-name
                        — no dscp-out-profile
                        — lsp-exp-in-profile lsp-exp-value
                        — no lsp-exp-in-profile
                        — lsp-exp-out-profile lsp-exp-value
                        — no lsp-exp-out-profile
                        — policer policier-id {queue queue-id}
                        — port-redirect-group {queue queue-id | policer policer-id [queue
                          queue-id]}
                        — no port-redirect-group
                — ingress
                        — default-action  fc fc-name profile {in | out}
                        — dot1p dot1p-priority fc fc-name profile {in | out | use-de}
                        — no dot1p dot1p-priority | use-de}
                        — no dot1p dot1p-priority
                        — dscp dscp-name fc fc-name  profile {in | out}
                        — no dscp dscp-name
                        — [no] fc fc-name
                                — fp-redirect-group policer policer-id
                                — no fp-redirect-group policer
                                — fp-redirect-group multicast-policer policer-id
                                — no fp-redirect-group-queue multicast-policer
                        — [no] ler-use-dscp
                        — lsp-exp lsp-exp-value fc fc-name profile {in | out}
                        — no lsp-exp lsp-exp-value
```

Self-Generated Traffic Commands

```
config
    — router
        — sgt-qos
            — application dscp-app-name dscp {dscp-value | dscp-name}
            — application dot1p-app-name dot1p dot1p-priority
            — no application {dscp-app-name | dot1p-app-name}
            — dscp dscp-name fc fc-name
            — no dscp dscp-name
```

Operational Commands

```
config
    — qos
        — copy network src-pol dst-pol [overwrite]
```

Show Commands

```
show
    — qos
        — dscp-table  value dscp-value
        — mc-fr-profile-ingress [detail]
        — mc-fr-profile-egress [detail]
```

— **network** *policy-id* [**detail**]**show**
— **router**
— **sgt-qos**
— **application** [*app-name*] [**dscp-dot1p**]
— **dscp-map** [*dscp-name*]

Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
no description

Context config>qos>network *policy-id*

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax **copy network** *src-pol dst-pol* [**overwrite**]

Context config>qos

Description This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id.

The **copy** command is used to create new policies using existing policies and also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters **network** *src-pol dst-pol* **—** Indicates that the source and destination policies are network policy IDs. Specify the source policy that the copy command will copy and specify the destination policy to which the command will duplicate the policy to a new or different policy ID.

Values 1 — 65535

overwrite — Specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, an error will occur if the destination policy ID exists.

```
SR>config>qos# copy network 1 427
MINOR: CLI Destination "427" exists use {overwrite}.
SR>config>qos# copy network 1 427 overwrite
```

scope

Syntax **scope** {**exclusive** | **template**}
no scope

Context config>qos>network *policy-id*

Description This command configures the network policy scope as exclusive or template. The policy's scope cannot be changed if the policy is applied to an interface.

The **no** form of this command sets the scope of the policy to the default of **template**.

Default template

Parameters **exclusive —** When the scope of a policy is defined as exclusive, the policy can only be applied to one interface. If a policy with an exclusive scope is assigned to a second interface an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface.
The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1.

template — When the scope of a policy is defined as template, the policy can be applied to multiple interfaces on the router.

Default QoS policies are configured with template scopes. An error is generated if you try to modify the template scope parameter to exclusive scope on default policies.

Multi-Link Frame Relay Commands

mc-fr-profile-ingress

Syntax [**no**] **mc-fr-profile-ingress** *profile-id*

Context config>qos

Description This command creates a profile for the user to configure the ingress QoS parameters of a Multi-Link Frame Relay (MLFR) bundle. A maximum of 128 ingress QoS profiles may be created on the system.

The **no** form of this command deletes the profile.

Default none

Parameters *profile-id* — Specifies the profile number.

Values 1 — 65535

class

Syntax **class** *class-id*

Context config>qos>mc-fr-profile-ingress
config>qos>mc-fr-profile-egress

Description This command provides the Frame Relay scheduling class context for the user to configure the ingress or egress QoS parameters of an MLFR bundle or an FRF.12 UNI/NNI link for this profile.

Default none

Parameters *class-id* — Specifies the Frame Relay scheduling class number.

Values 0 — 3

reassembly-timeout

Syntax **reassembly-timeout** *timeout-value*
no reassembly-timeout

Context config>qos>mc-fr-profile-ingress>class

Description This command configures the value of the MLFR bundle ingress per-class reassembly timer for the profile.

Default Class 0=10 msec
Class 1=10 msec
Class 2=100 msec
Class 3=1000 msec

Parameters *timeout-value* — Specifies the timeout value, in milliseconds.

Values 1 — 1000

mc-fr-profile-egress

Syntax [**no**] **mc-fr-profile-egress** *profile-id*

Context config>qos

Description This command creates a profile for the user to configure the egress QoS parameters of an MLFR bundle or an FRF.12 UNI/NNI link. A maximum of 128 egress QoS egress profile may be created on the system.

The no form of this command deletes the profile.

Default none

Parameters *profile-id* — Specifies the profile number.

Values 1 — 65535

max-queue-size

Syntax **max-queue-size** *queue-size*
no max-queue-size

Context config>qos>mc-fr-profile-egress>class

Description This command configures the maximum size for each Frame Relay scheduling class queue for this profile.

Default Class 0=10
Class 1=50
Class 2=150
Class 3=750

Parameters *queue-size* — Specifies the number, in milliseconds, of the available link or bundle rate.

Values 1 — 1000

mir

Syntax **mir** *mir*
no mir

Context config>qos>mc-fr-profile-egress>class

Description This command configures the minimum information rate scheduling parameter for each Frame Relay scheduling class queues for this profile.

Default 90% for all classes

Parameters *mir* — Specifies the percentage of the available link or bundle rate.

Values 1 — 100

weight

Syntax **weight** *weight*
no weight

Context config>qos>mc-fr-profile-egress>class

Description This command configures the WRR weight scheduling parameter for each Frame Relay scheduling class queue for this profile.

Default Class 0=N/A

Class 1=1 (not configurable)

Class 2=89

Class 3=10

Parameters *weight* — Specifies the weight schedule.

Values 1 — 100

Network QoS Policy Commands

network

Syntax **network** *network-policy-id* [**create**]
no network *network-policy-id*

Context config>qos

Description This command creates or edits a QoS network policy. The network policy defines the treatment IP or MPLS packets receive as they ingress and egress the network port.

The QoS network policy consists of an ingress and egress component. The ingress component of the policy defines how DiffServ code points and MPLS EXP bits are mapped to internal forwarding class and profile state. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the router. The mapping on each network interface defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the network interface.

The egress component of the network QoS policy defines the queuing parameters associated with each forwarding class. Each of the forwarding classes defined within the system automatically creates a queue on each network interface. This queue gets all the parameters defined within the default network QoS policy 1 until an explicit policy is defined for the network interfaceaccess uplink port. If the egressing packet originated on an ingress SAP, or the remarking parameter is defined for the egress interface, the egress QoS policy also defines the IP DSCP or MPLS EXP bit marking based on the forwarding class and the profile state.

Network **policy-id 1** exists as the default policy that is applied to all network interfaces by default. The network **policy-id 1** cannot be modified or deleted. It defines the default DSCP-to-FC mapping and MPLS EXP-to-FC mapping and for the ingress. For the egress, it defines six forwarding classes which represent individual queues and the packet marking criteria.

Network policy-id 1 exists as the default policy that is applied to all network ports by default. This default policy cannot be modified or deleted. It defined the default DSCP-to-FC mapping and default unicast meters for ingress IP traffic. For the egress, if defines the forwarding class to Dot1p and DSCP values and the packet marking criteria.

If a new network policy is created (for instance, policy-id 3), only the default action and egress forwarding class parameters are identical to the default policy. A new network policy does not contain the default DSCP-to-FC and MPLS-EXP-to-FC mapping for network QoS policy of type **ip-interface** or the DSCP-to-FC mapping (for network QoSpolicy of type **port**). The default network policy can be copied (use the copy command) to create a new network policy that includes the default ingress DSCP-to-FC and MPLS EXP-to-FC mapping (as appropriate). You can modify parameters or use the **no** modifier to remove an object from the configuration.

Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all network interfaces where this policy is applied. For this reason, when many changes are required on a policy, it is highly recommended that the policy be copied to a work area policy-id. That work-in-progress policy can be modified until complete and then written over the original policy-id. Use the config qos copy command to maintain policies in this manner.

The **no** form of this command deletes the network policy. A policy cannot be deleted until it is removed from all entities where it is applied. The default network **policy** *policy-id* 1 cannot be deleted.

Default System Default Network Policy 1

Parameters *network-policy-id* — The policy-id uniquely identifies the policy on the router.

Default none

Values 1— 65535

Network Ingress QoS Policy Commands

ingress

Syntax **ingress**

Context config>qos>network *policy-id*

Description This command is used to enter the CLI node that creates or edits policy entries that specify the DiffServ code points to forwarding class mapping for all IP packets and define the MPLS EXP bits to forwarding class mapping for all labeled packets.

When pre-marked IP or MPLS packets ingress on a network port, they get a Per Hop Behavior (that is, the QoS treatment through the router-based on the mapping defined under the current node.

default-action

Syntax **default-action fc** *fc-name* **profile** {**in** | **out**}

Context config>qos>network>ingress

Description This command defines or edits the default action to be taken for packets that have an undefined DSCP or MPLS EXP bits set. The **default-action** command specifies the forwarding class to which such packets are assigned.

Multiple default-action commands will overwrite each previous default-action command.

Default default-action fc be profile out

Parameters **fc** *fc-name* **—** Specify the forwarding class name. All packets with DSCP value or MPLS EXPor dot1p bits bits that is not defined will be placed in this forwarding class.

Default None, the fc name must be specified

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** All packets that are assigned to this forwarding class will be considered in or out of profile based on this command. In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets.

Default None

Values in, out

dot1p

Syntax **dot1p** *dot1p-priority* **fc** *fc-name* **profile** {**in** | **out | use-de**}
no dot1p *dot1p-priority*

Context config>qos>network>ingress

Description This command explicitly sets the forwarding class or enqueuing priorityand profile of the packet when a packet is marked with a *dot1p-priority* specified. Adding a dot1p rule on the policy forces packets that match the *dot1p-priority* specified to override be assigned to the forwarding class and enqueuing priority and profile of the packet based on the parameters included in the Dot1p rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueuing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

The *dot1p-priority* is derived from the most significant three bits in the IEEE 802.1Q or IEEE 802.1P header. The three dot1p bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop Quality-of-Service (QoS) behavior.

The **no** form of this command removes the explicit dot1p classification rule from the policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

Parameters *dot1p-priority* — This value is a required parameter that specifies the unique IEEE 802.1P value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

Values 0 — 7

fc *fc-name* **—** The value given for the *fc-name* parameter must be one of the predefined forwarding classes in the system. Specifying the *fc-name* is optional. When a packet matches the rule, the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

Values `be`, `l2`, `af`, `l1`, `h2`, `ef`, `h1`, `nc`

profile {**in** | **out | use-de**} **—** All packets that are assigned to this forwarding class will be considered in or out of profile based on this command or to use the DE1 bit to determine the profile of the packets. In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets.

Default none, the profile name must be specified.

dscp

Syntax **dscp** *dscp-name* **fc** *fc-name* **profile** {**in** | **out**}
no dscp *dscp-name*

Context config>qos>network *policy-id*>ingress

Description This command creates a mapping between the DiffServ Code Point (DSCP) of the network ingress traffic and the forwarding class.

Ingress traffic that matches the specified DSCP will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all sixty-four DiffServ code points to the forwarding class. For undefined code points, packets are assigned to the forwarding class specified under the **default-action** command.

The **no** form of this command removes the DiffServ code point to forwarding class association. The **default-action** then applies to that code point value.

Default none

Parameters *dscp-name* — The name of the DiffServ code point to be associated with the forwarding class. DiffServ code point can only be specified by its name and only an existing DiffServ code point can be specified. The software provides names for the well known code points.

The system-defined names available are as follows. The system-defined names must be referenced as all lower case exactly as shown in the first column in Table 24 and Table 25 below.

Additional names to code point value associations can be added using the '**dscp-name** *dscp-name dscp-value*' command.

The actual mapping is being done on the *dscp-value*, not the *dscp-name* that references the *dscp-value*. If a second *dscp-name* that references the same *dscp-value* is mapped within the policy, an error will occur. The second name will not be accepted until the first name is removed.

Table 24: Default DSCP Names to DSCP Value Mapping Table

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary
nc1	48	0x30	0b110000
nc2	56	0x38	0b111000
ef	46	0x2e	0b101110
af41	34	0x22	0b100010
af42	36	0x24	0b100100
af43	38	0x26	0b100110
af31	26	0x1a	0b011010
af32	28	0x1c	0b011100
af33	30	0x1d	0b011110
af21	18	0x12	0b010010
af22	20	0x14	0b010100
af23	22	0x16	0b010110
af11	10	0x0a	0b001010
af12	12	0x0c	0b001100
af13	14	0x0e	0b001110
default	0	0x00	0b000000

Table 25: Default Class Selector Code Points to DSCP Value Mapping Table

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary
cs7	56	0x38	0b111000
cs6	48	0X30	0b110000
cs5	40	0x28	0b101000
cs4	32	0x20	0b100000

Table 25: Default Class Selector Code Points to DSCP Value Mapping Table (Continued)

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary
cs3	24	0x18	0b011000
cs2	16	0x10	0b010000
cs1	08	0x8	0b001000

fc *fc-name* **—** Enter this required parameter to specify the *fc-name* with which the code point will be associated.

Default none, for every DSCP value defined, the forwarding class must be indicated.

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** Enter this required parameter to indicate whether the DiffServ code point value is the in-profile or out-of-profile value.

NOTE 1: DSCP values mapping to forwarding classes Expedited (ef), High-1 (h1) and Network-Control (nc) can only be set to in-profile.

NOTE 2: DSCP values mapping to forwarding class 'be' can only be set to out-of-profile.

Default None, for every DSCP value defined, the profile must be indicated. If a DSCP value is not mapped, the default-action forwarding class and profile state will be used for that value.

Values in, out

fp-redirect-group

Syntax **fp-redirect-group policer** *policer-id*
no fp-redirect-group policer

Context config>qos>network>ingress>fc

Description This command is used to redirect the FC of a packet of a PW or network IP interface to an ingress forwarding plane queue-group.

It defines the mapping of a FC to a policer-id and redirects the lookup of the policer of the same ID in some ingress forwarding plane queue-group instance. However, the queue-group name and instance are explicitly provided only at the time the network QoS policy is applied to the ingress context of a spoke-sdp or a network IP interface.

The **no** version of this command removes the redirection of the FC.

Parameters **policer** *policer-id* — The specified *policer-id* must exist within the queue-group template applied to the ingress context of the forwarding plane.

Values 1—8

fp-redirect-group

Syntax **fp-redirect-group multicast-policer** *policer-id*
no fp-redirect-group multicast-policer

Context config>qos>network>ingress>fc

Description This command is used to redirect the FC of a multicast packet of a PW or network IP interface to an ingress forwarding plane queue-group.

It defines the mapping of a FC to a policer-id and redirects the lookup of the policer of the same ID in some ingress forwarding plane queue-group instance. However, the queue-group name and instance are explicitly provided only at the time the network QoS policy is applied to the ingress context of a spoke-sdp or a network IP interface.

The **no** version of this command removes the redirection of the FC.

Parameters **multicast** *policer-id* — The specified *policer-id* must exist within the queue-group template applied to the ingress context of the forwarding plane.

Values 1—32

ler-use-dscp

Syntax [**no**] **ler-use-dscp**

Context config>qos>network>ingress

Description This command is used to enable tunnel QoS mapping on all ingress network IP interfaces the network-qos-policy-id is associated with. The command may be defined at anytime after the network QoS policy has been created. Any network IP interfaces currently associated with the policy will immediately start to use the internal IP ToS field of any tunnel terminated IP routed packet received on the interface, ignoring any QoS markings in the tunnel portion of the packet.

This attribute provides the ability to ignore the network ingress QoS mapping of a terminated tunnel containing an IP packet that is to be routed to a base router or VPRN destination. This is advantageous when the mapping for the tunnel QoS marking does not accurately or completely reflect the required QoS handling for the IP routed packet. When the mechanism is enabled on an ingress network IP interface, the IP interface will ignore the tunnel's QoS mapping and derive the internal forwarding class and profile based on the precedence or DiffServe Code Point (DSCP) values within the routed IP header ToS field compared to the Network QoS policy defined on the IP interface.

The default state is not to enforce tunnel termination IP routed QoS override within the network QoS policy.

The **no** form of the command removes tunnel termination IP routed QoS override from the network QoS policy and all ingress network IP interfaces associated with the policy.

Default no ler-use-dscp

lsp-exp

Syntax **lsp-exp** *lsp-exp-value* **fc** *fc-name* **profile** {**in** | **out**}
no lsp-exp *lsp-exp-value*

Context config>qos>network *policy-id*>ingress

Description This command creates a mapping between the LSP EXP bits of the network ingress traffic and the forwarding class.

Ingress traffic that matches the specified LSP EXP bits will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all eight LSP EXP bit values to the forwarding class. For undefined values, packets are assigned to the forwarding class specified under the **default-action** command.

The **no** form of this command removes the association of the LSP EXP bit value to the forwarding class. The **default-action** then applies to that LSP EXP bit pattern.

Default none

Parameters *lsp-exp-value* — Specify the LSP EXP values to be associated with the forwarding class.

Default None, the lsp-exp command must define a value.

Values 0 to 8 (Decimal representation of three EXP bit field)

fc *fc-name* **—** Enter this required parameter to specify the fc-name that the EXP bit pattern will be associated with.

Default None, the lsp-exp command must define a fc-name.

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** Enter this required parameter to indicate whether the LSP EXP value is the in-profile or out-of-profile value.

Default None, the lsp-exp command must define a profile state.

Values in, out

Network Egress QoS Policy Commands

egress

Syntax **egress**

Context config>qos>network *policy-id*

Description This command is used to enter the CLI node that creates or edits egress policy entries that specify the forwarding class queues to be instantiated when this policy is applied to the network port.

The forwarding class and profile state mapping to in and out-of-profile DiffServ code points and MPLS EXP bits mapping for all labeled packets are also defined in this context.

All service packets are aggregated into DiffServ based egress queues on the network interface. The service packets are transported either with IP GRE encapsulation or over a MPLS LSP. The exception is with the IES service. In this case, the actual customer IP header has the DSCP field mapped.

All out-of-profile service packets are marked with the corresponding out-of-profile DSCP or the EXP bit value at network egress. All the in-profile service ingress packets are marked with the corresponding in-profile DSCP or EXP bit value based on the forwarding class they belong.

fc

Syntax [**no**] **fc** *fc-name*

Context config>qos>network>egress

Description This command specifies the forwarding class name. The forwarding class name represents an egress queue. The **fc** *fc-name* represents a CLI parent node that contains sub-commands or parameters describing the egress characteristics of the queue and the marking criteria of packets flowing through it. The **fc** command overrides the default parameters for that forwarding class to the values defined in the network default policy.

The **no** form of this command removes the forwarding class name associated with this queue, as appropriate. The forwarding class reverts to the defined parameters in the default network policy. If the *fc-name* is removed from the network policy that forwarding class reverts to the factory defaults.

Default Undefined forwarding classes default to the configured parameters in the default network policy policy-id 1.

Parameters *fc-name* — The case-sensitive, system-defined forwarding class name for which policy entries will be created.

Default none

Values be, l2, af, l1, h2, ef, h1, nc

Network Egress QoS Policy Forwarding Class Commands

de-mark

Syntax **de-mark** [**force** *de-value*]
no de-mark

Context config>qos>network>egress>fc

Description This command is used to explicitly define the marking of the DE bit for **fc** *fc-name* according to the in and out of profile status of the packet (fc-name may be used to identify the dot1p-value).

If no de-value is present, the default values are used for the marking of the DE bit: i.e. 0 for in-profile packets, 1 for out-of-profile ones – see 802.1ad-2005 standard.

In the PBB case, for a Network Port (B-SDP), the following rules must be used:

- the outer VID follows the rules for regular SDP
- for packets originated from a local I-VPLS/PBB-Epipe, this command dictates the marking of the DE bit for both the outer (link level) BVID and ITAG; if the command is not used the DE bit will be set to zero.
- for transit packets (B-SAP/B-SDP to B-SDP) the related ITAG bits will be preserved, same for BVID.

If the de-value is specifically mentioned in the command line it means this value is to be used for all the packets of this forwarding class regardless of their in/out of profile status.

Values 0 or 1

dot1p

Syntax **dot1p** *dot1p-priority*
no dot1p

Context config>qos>network>egress>fc

Description This command will be used whenever the dot1p bits are set to a common value regardless of the internal in | out-profile of the packets. Although it is not mandatory, it is expected that this command is used in combination with the de-mark command to enable the marking of the DE bit according to the internal profile of the packet.

This command acts as a shortcut version of configuring the two existing commands with the same dot1p-priority.

To minimize the required changes the dot1p x command should be saved in the configuration as dot1p-in-profile x and dot1p-out-profile x.

dot1p-in-profile

Syntax **dot1p-in-profile** *dot1p-priority*
no dot1p-in-profile

Context config>qos>network>egress>fc *fc-name*

Description This command specifies dot1p in-profile mappings.

The **no** form of the command reverts to the default in-profile *dot1p-priority* setting for policy-id 1.

Parameters *dot1p-priority* — This value is a required parameter that specifies the unique IEEE 802.1P value that will match the Dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

Values 0 — 7

dot1p-out-profile

Syntax **dot1p-out-profile** *dot1p-priority*
no dot1p-out-profile

Context config>qos>network>egress>fc *fc-name*

Description This command specifies dot1p out-profile mappings.

The **no** form of the command reverts to the default out-profile *dot1p-priority* setting for policy-id 1.

Parameters *dot1p-priority* — This value is a required parameter that specifies the unique IEEE 802.1P value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

Values 0 — 7

dscp-in-profile

Syntax **dscp-in-profile** *dscp-name*
no dscp-in-profile

Context config>qos>network *policy-id*>egress>fc *fc-name*

Description This command specifies the in-profile DSCP name for the forwarding class. The corresponding DSCP value will be used for all IP packets requiring marking the egress on this forwarding class queue that are in profile.

When multiple DSCP names are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.

The **no** form of this command reverts to the factory default in-profile dscp-name setting for policy-id 1.

Default Policy-id 1: Factory setting

Policy-id 2 — 65535: Policy-id 1 setting

Parameters *dscp-name* — System- or user-defined, case-sensitive *dscp-name*.

Default none

Values Any defined system- or user-defined *dscp-name*

dscp-out-profile

Syntax **dscp-out-profile** *dscp-name*
no dscp-out-profile

Context config>qos>network *policy-id*>egress>fc *fc-name*

Description This command specifies the out-of-profile DSCP name for the forwarding class. The corresponding DSCP value will be used for all IP packets requiring marking the egress on this forwarding class queue that are out-of-profile.

When multiple DSCP names are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.

The **no** form of this command reverts to the factory default out-of-profile dscp-name setting for policy-id 1.

Default Policy-id 1: Factory setting

Policy-id 2 — 65535: Policy-id 1 setting

Parameters *dscp-name* — System- or user-defined, case-sensitive *dscp-name*.

Default none

Values Any defined system- or user-defined *dscp-name*

lsp-exp-in-profile

Syntax **lsp-exp-in-profile** *lsp-exp-value*
no lsp-exp-in-profile

Context config>qos>network *policy-id*>egress>fc *fc-name*

Description This command specifies the in-profile LSP EXP value for the forwarding class. The EXP value will be used for all LSP labeled packets requiring marking the egress on this forwarding class queue that are in-profile.

When multiple EXP values are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.

The **no** form of this command reverts to the factory default in-profile EXP setting.

Default Policy-id 1: Factory setting

Policy-id 2 — 65535: Policy-id setting

Parameters *lsp-exp-value* — The 3-bit LSP EXP bit value, expressed as a decimal integer.

Default none

Values 0 — 7

lsp-exp-out-profile

Syntax **lsp-exp-out-profile** *lsp-exp-value*
no lsp-exp-out-profile

Context config>qos>network *policy-id*>egress>fc *fc-name*

Description This command specifies the out-of-profile LSP EXP value for the forwarding class. The EXP value will be used for all LSP labeled packets requiring marking the egress on this forwarding class queue that are out-of-profile.

When multiple EXP values are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.

The **no** form of this command reverts to the factory default out-of-profile EXP setting.

Default Policy-id 1: Factory setting

Policy-id 2 — 65535: Policy-id setting

Parameters *mpls-exp-value* — The 3-bit MPLS EXP bit value, expressed as a decimal integer.

Default none

Values 0 — 7

policer

Syntax **policer** *policer-id*
no policer

Context *config>qos>queue-group-templates>ingress>queue-group*
config>qos>queue-group-templates>egress>queue-group

Description This command is used in ingress and egress queue-group templates to create, modify, or delete a policer.

Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The ingress queue-group template have up to 32 policers (numbered 1 through 32) and can be defined while the egress queue-group template supports a maximum of 8 (numbered 1 through 8). While a policer can be defined in a queue-group template, it is not actually created until the queue-group template is instantiated on ingress context of a forwarding plane or on the egress context of a port.

Once a policer is created, the policer's metering rate and profiling rates can be defined as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues which have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

Once a policer is created, it cannot be deleted from the queue-group template unless any forwarding classes that are redirected to the policer are first removed.

The **no** version of this command deletes the policer.

Parameters *policer-id* — The policer-id must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer’s parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword requirements which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer’s context for possible parameter modification..

Values 1—32 ingress

Values 1—8 egress

port-redirect-group

Syntax **port-redirect-group** {**queue** *queue-id* | **policer** *policer-id* [**queue** *queue-id*]}
no port-redirect-group

Context config>qos>network>egress>fc

Description This command is used to redirect the FC of a packet of a PW or network IP interface to an egress port queue-group.

It defines the mapping of a FC to a queue-id or a policer-id and a queue-id, and redirects the lookup of the queue or policer of the same ID in some egress port queue-group instance. However, the queue-group name and instance are explicitly provided only at the time the network QoS policy is applied to egress context of a spoke-sdp or a network IP interface.

The **no** version of this command removes the redirection of the FC.

Parameters *queue-id* — This parameter must be specified when executing the **port-redirect-group** command. The specified *queue-id* must exist within the egress port queue group on each IP interface where the network QoS policy is applied.

Values 1 — 8

policer id — The specified policer-id must exist within the queue-group template applied to the ingress context of the forwarding plane.

Values 1 — 8

remarking

Syntax [**no**] **remarking** [**force**]

Context config>qos>network *policy-id*>egress

Description This command remarks both customer traffic and egress network IP interface traffic; VPRN customer traffic is not remarked. The remarking is based on the forwarding class to DSCP and LSP EXP bit mapping defined under the egress node of the network QoS policy.

Normally, packets that ingress on network ports have either DSCP or, in case of MPLS packets, LSP EXP bit set by an upstream router. The packets are placed in the appropriate forwarding class based on the DSCP to forwarding class mapping or the LSP EXP to forwarding class mapping. The DSCP or LSP EXP bits of such packets are not altered as the packets egress this router, unless **remarking** is enabled.

Remarking can be required if this router is connected to a different DiffServ domain where the DSCP to forwarding class mapping is different.

Normally no remarking is necessary when all router devices are in the same DiffServ domain.

The network QoS policy supports an egress flag that forces remarking of packets that were received on trusted IES and network IP interfaces. This provides the capability of remarking without regard to the ingress state of the IP interface on which a packet was received. The effect of the setting of the egress network remark trusted state on each type of ingress IP interface and trust state is shown in the following table.

The remark trusted state has no effect on packets received on an ingress VPRN IP interface.

Ingress IP Interface Type and Trust State	Egress Network IP Interface Trust Remark Disabled (Default)	Egress Network IP Interface Trust Remark Enabled
IES Non-Trusted (Default)	Egress Remarked	Egress Remarked
IES Trusted	Egress Not Remarked	Egress Remarked
VPRN Non-Trusted	Egress Remarked	Egress Remarked
VPRN Trusted (Default)	Egress Not Remarked	Egress Not Remarked
Network Non-Trusted	Egress Remarked	Egress Remarked
Network Trusted (Default)	Egress Not Remarked	Egress Remarked

The **no** form of this command reverts to the default behavior.

Default **no remarking** — Remarking disabled in the Network QoS policy.

Parameters **force —** Specifies that all IP routed traffic egressing the associated network interface will have its EXP, DSCP, P-bit and DE bit setting remarked as defined in the associated QoS policy. Only bit fields configured in the QoS policy will be remarked; all others will be left untouched or set based on the default if the fields were not present at ingress.

Self-Generated Traffic Commands

sgt-qos

Syntax **sgt-qos**

Context config>router

Description This command enables the context to configure DSCP/Dot1p re-marking for self-generated traffic.

application

Syntax **application** *dscp-app-name* **dscp** {*dscp-value* | *dscp-name*}
application *dot1p-app-name* **dot1p** *dot1p-priority*
no application {*dscp-app-name* | *dot1p-app-name*}

Context config>router>sgt-qos

Description This command configures DSCP/Dot1p re-marking for self-generated application traffic. When an application is configured using this command, then the specified DSCP name/value is used for all packets generated by this application within the router instance it is configured. The instances can be base router, vprn or management.

Using the value configured in this command:

- Sets the DSCP bits in the IP packet.
- Maps to the FC. This value will be signaled from the CPM to the egress forwarding complex.
- Based on this signaled FC the egress forwarding complex QoS policy sets the IEEE802.1 dot1P and LSP EXP bits.
- The Dot1P and the LSP EXP bits are set by the egress complex for all packets based on the signaled FC. This includes ARP, PPPoE and IS-IS packets that, due to their nature, do not carry DSCP bits.
- The DSCP value in the egress IP header will be as configured in this command. The egress QoS policy will not overwrite this value.

Only one DSCP name/value can be configured per application, if multiple entries are configured then the subsequent entry overrides the previous configured entry.

The **no** form of this command reverts back to the default value.

Parameters *dscp-app-name* — Specifies the DSCP application name.

Values ldp, rsvp, bgp, rip, msdp, pim, ospf, igmp, mld, telnet, tftp, ftp, ssh, snmp, snmp-notification, syslog, icmp, traceroute, tacplus, dns, ntp, radius, cflowd, dhcp, ndis, vrrp, srrp

dscp-value — Specifies a value when this packet egresses the respective egress policy should provide the mapping for the DSCP value to either LSP-EXP bits or IEEE 802.1p (Dot1P) bits as appropriate otherwise the default mapping applies.

Values 0 — 63

dscp-name — Specifies the DSCP name.

Values none, be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

dot1p-priority — Specifies the Dot1P priority.

Values 0 — 7

dot1p-app-name — Specifies the Dot1P application name.

Values arp, isis, pppoe

dscp

Syntax **dscp** *dscp-name* **fc** *fc-name*
no dscp *dscp-name*

Context config>router>sgt-qos

Description This command creates a mapping between the DiffServ Code Point (DSCP) of the self generated traffic and the forwarding class.

Self generated traffic that matches the specified DSCP will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all sixty-four DiffServ code points to the forwarding class.

All dscp name that defines a dscp value must be explicitly defined

The **no** form of this command removes the DiffServ code point to forwarding class association.

Default none

Parameters *dscp-name* — The name of the DiffServ code point to be associated with the forwarding class. DiffServ code point can only be specified by its name and only an existing DiffServ code point can be specified. The software provides names for the well known code points.

Values be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44,

cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

fc *fc-name* **—** Specify the forwarding class name. All packets with DSCP value or MPLS EXP bits that is not defined will be placed in this forwarding class.

Default None, the fc name must be specified

Values be, l2, af, l1, h2, ef, h1, nc

Show Commands

dscp-table

Syntax **dscp-table** [**value** *dscp-value*]

Context show>qos

Description Displays the DSCP name to DSCP value mappings.

Parameters **value** *dscp-value* **—** The specific DSCP value for which to display information.

Default Show all values

Values 0 — 63

Table 26: Show QoS Network Table Output Fields

Label	Description
DSCP Name	Displays the name of the DiffServ code point to be associated with the forwarding class.
DSCP Value	Displays the DSCP values range between 0 and 63.
TOS (bin)	Displays the type of service in Binary format.
TOS (hex)	Displays the type of service in Hex format.

Sample Output

```
A:ALA-48# show qos dscp-table
===========================================================
DSCP Mapping
===========================================================
DSCP Name      DSCP Value     TOS (bin)      TOS (hex)
-----------------------------------------------------------
be             0              0000 0000      00
cp1            1              0000 0100      04
cp2            2              0000 1000      08
cp3            3              0000 1100      0C
cp4            4              0001 0000      10
cp5            5              0001 0100      14
cp6            6              0001 1000      18
cp7            7              0001 1100      1C
cs1            8              0010 0000      20
cp9            9              0010 0100      24
af11           10             0010 1000      28
cp11           11             0010 1100      2C
```

```
af12            12              0011 0000       30
cp13            13              0011 0100       34
af13            14              0011 1000       38
cp15            15              0011 1100       3C
cs2             16              0100 0000       40
cp17            17              0100 0100       44
af21            18              0100 1000       48
cp19            19              0100 1100       4C
af22            20              0101 0000       50
cp21            21              0101 0100       54
af23            22              0101 1000       58
cp23            23              0101 1100       5C
cs3             24              0110 0000       60
cp25            25              0110 0100       64
af31            26              0110 1000       68
cp27            27              0110 1100       6C
af32            28              0111 0000       70
cp29            29              0111 0100       74
af33            30              0111 1000       78
cp31            31              0111 1100       7C
cs4             32              1000 0000       80
cp33            33              1000 0100       84
af41            34              1000 1000       88
cp35            35              1000 1100       8C
af42            36              1001 0000       90
cp37            37              1001 0100       94
af43            38              1001 1000       98
cp39            39              1001 1100       9C
cs5             40              1010 0000       A0
cp41            41              1010 0100       A4
cp42            42              1010 1000       A8
cp43            43              1010 1100       AC
cp44            44              1011 0000       B0
cp45            45              1011 0100       B4
ef              46              1011 1000       B8
cp47            47              1011 1100       BC
nc1             48              1100 0000       C0
cp49            49              1100 0100       C4
cp50            50              1100 1000       C8
cp51            51              1100 1100       CC
cp52            52              1101 0000       D0
cp53            53              1101 0100       D4
cp54            54              1101 1000       D8
cp55            55              1101 1100       DC
nc2             56              1110 0000       E0
cp57            57              1110 0100       E4
cp58            58              1110 1000       E8
cp59            59              1110 1100       EC
cp60            60              1111 0000       F0
cp61            61              1111 0100       F4
cp62            62              1111 1000       F8
cp63            63              1111 1100       FC
===============================================================
A:ALA-48#


A:ALA-48# show qos dscp-table value 46
===============================================================
DSCP Mapping
```

```
===========================================================
DSCP Name      DSCP Value     TOS (bin)      TOS (hex)
-----------------------------------------------------------
ef             46             1011 1000      B8
===========================================================
A:ALA-48#
```

mc-fr-profile-ingress

Syntax **mc-fr-profile-ingress** [**detail**]

Context show>qos

Description This command displays MLFR ingress profile details.

Sample Output

```
*A:Cpm-A# show qos mc-fr-profile-ingress
===============================================================================
Multi-class Frame-Relay Ingress Profiles
===============================================================================
Profile-Id  Description
-------------------------------------------------------------------------------
1           Default ingress multi-class frame-relay profile.
===============================================================================
*A:Cpm-A#
*A:Cpm-A# show qos mc-fr-profile-ingress 1 detail
===============================================================================
Multi-class FR Ingress Profile (1)
===============================================================================
Profile-id : 1
Description: Default ingress multi-class frame-relay profile.
-------------------------------------------------------------------------------
FR Class     Reassembly Timeout
-------------------------------------------------------------------------------
0            10
1            10
2            100
3            1000
===============================================================================
Associations
-------------------------------------------------------------------------------
No Matching Entries
```

mc-fr-profile-egress

Syntax **mc-fr-profile-egress** [**detail**]

Context show>qos

Description This command displays MLFR egress profile details.

Sample Output

```
*A:Cpm-A# show qos mc-fr-profile-egress 1
===============================================================================
Multi-class FR Egress Profile (1)
===============================================================================
Profile-id : 1
Description: Default egress multi-class frame-relay profile.
===============================================================================
*A:Cpm-A#
*A:Cpm-A# show qos mc-fr-profile-egress 1 detail
===============================================================================
Multi-class FR Egress Profile (1)
===============================================================================
Profile-id : 1
Description: Default egress multi-class frame-relay profile.
===============================================================================
MCFR        Mir         Weight      Max Size
Class
-------------------------------------------------------------------------------
0           100         0           25
1           85          0           5
2           0           66          200
3           0           33          1000
===============================================================================
Associations
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:Cpm-A#
```

network

Syntax **network** [*policy-id*] [**detail**]

Context show>qos

Description This command displays network policy information.

Parameters *policy-id* — Displays information for the specific policy ID.

Default all network policies

Values 1 — 65535

detail — Includes information about ingress and egress DSCP and LSP EXP bit mappings and network policy interface associations.

Network QoS Policy Output Fields — The following table describes network QoS Policy output fields.

Table 27: Show QoS Network Output Fields

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Remark	True – Remarking is enabled for all packets that egress this router where the network policy is applied. The remarking is based on the forwarding class to DSCP and LSP EXP bit mapping defined under the egress node of the network QoS policy. False – Remarking is disabled.
Description	A text string that helps identify the policy's context in the configuration file.
Forward Class/ FC Name	Specifies the forwarding class name.
Profile	Out – Specifies that IP packets requiring marking the egress on this forwarding class queue that are out of profile. In – Specifies that IP packets requiring marking the egress on this forwarding class queue that are in profile.
Accounting	Packet-based – Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow. Frame-based – Specifies that the meters associated with this policy account for the packet framing overheads (such as for Ethernet the IFG and preamble), while accounting the bandwidth to be used by the flow.
DSCP Mapping:	
Out-of-Profile	Displays the DSCP used for out-of-profile traffic.
In-Profile	Displays the DSCP used for in-profile traffic.
LSP EXP Bit Mapping:	
Out-of-Profile	Displays the LSP EXP value used for out-of-profile traffic.
In-Profile	Displays the LSP EXP value used for in-profile traffic.
Interface	Displays the interface name.

Table 27: Show QoS Network Output Fields (Continued)

Label	Description
IP Addr	Displays the interface IP address.
Port-Id	Specifies the physical port identifier that associates the interface.

```
A:ALA-12# show qos network
===============================================================================
Network Policies
===============================================================================
Policy-Id            Remark    Description
-------------------------------------------------------------------------------
1                    True Default network QoS policy.
===============================================================================
A:ALA-12#

A:ALA-12# show qos network 1
===============================================================================
QoS Network Policy
===============================================================================
Network Policy (1)
-------------------------------------------------------------------------------
Policy-id      : 1                           Remark       : True
Forward Class  : be                          Profile      : Out-profile
Description    : Default network QoS policy.
===============================================================================
A:ALA-12#



A:ALA-12# show qos network 1 detail
===============================================================================
QoS Network Policy
===============================================================================
Network Policy (1)
-------------------------------------------------------------------------------
Policy-id      : 1                           Remark       : True
Forward Class  : be                          Profile      : Out-profile
Description    : Default network QoS policy.

-------------------------------------------------------------------------------
DSCP                                  Fowarding Class                Profile
-------------------------------------------------------------------------------
ef                                    ef                              In
nc1                                   h1                              In
nc2                                   nc                              In
af11                                  af                              In
af12                                  af                              Out
af13                                  af                              Out
af21                                  l1                              In
af22                                  l1                              Out
af23                                  l1                              Out
af31                                  l1                              In
af32                                  l1                              Out
af33                                  l1                              Out
```

```
af41                                    h2                              In
af42                                    h2                              Out
af43                                    h2                              Out

-------------------------------------------------------------------------------
LSP EXP Bit Map                         Fowarding Class                 Profile
-------------------------------------------------------------------------------
0                                       be                              Out
1                                       l2                              In
2                                       af                              Out
3                                       af                              In
4                                       h2                              In
5                                       ef                              In
6                                       h1                              In
7                                       nc                              In

-------------------------------------------------------------------------------
Egress Forwarding Class Queuing
-------------------------------------------------------------------------------
FC Name        : af
- DSCP Mapping
Out-of-Profile : af12                          In-Profile   : af11
- LSP EXP Bit Mapping
Out-of-Profile : 2                             In-Profile   : 3

FC Name        : be
- DSCP Mapping
Out-of-Profile : default                       In-Profile   : default
- LSP EXP Bit Mapping
Out-of-Profile : 0                             In-Profile   : 0

FC Name        : ef
- DSCP Mapping
Out-of-Profile : ef                            In-Profile   : ef
- LSP EXP Bit Mapping
Out-of-Profile : 5                             In-Profile   : 5

FC Name        : h1
- DSCP Mapping
Out-of-Profile : nc1                           In-Profile   : nc1
- LSP EXP Bit Mapping
Out-of-Profile : 6                             In-Profile   : 6

FC Name        : h2
- DSCP Mapping
Out-of-Profile : af42                          In-Profile   : af41
- LSP EXP Bit Mapping
Out-of-Profile : 4                             In-Profile   : 4

FC Name        : l1
- DSCP Mapping
Out-of-Profile : af22                          In-Profile   : af21
- LSP EXP Bit Mapping
Out-of-Profile : 2                             In-Profile   : 3

FC Name        : l2
- DSCP Mapping
Out-of-Profile : cs1                           In-Profile   : cs1
- LSP EXP Bit Mapping
```

```
Out-of-Profile : 1                              In-Profile   : 1

FC Name        : nc
- DSCP Mapping
Out-of-Profile : nc2                            In-Profile   : nc2
- LSP EXP Bit Mapping
Out-of-Profile : 7                              In-Profile   : 7
-------------------------------------------------------------------------------
Interface Association
-------------------------------------------------------------------------------
Interface      : system
IP Addr.       : 10.10.0.3/32                   Port Id      : vport-1
Interface      : to-ser1
IP Addr.       : 10.10.13.3/24                  Port Id      : 1/1/2
===============================================================================
A:ALA-12#


config>qos# show qos network 2 detail
===============================================================================
QoS Network Policy
-------------------------------------------------------------------------------
Network Policy (2)
-------------------------------------------------------------------------------
Policy-id      : 2                              Remark       : True
Forward Class  : be                             Profile      : Out
LER Use DSCP   : False
-------------------------------------------------------------------------------
DSCP            Forwarding Class    Profile
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
LSP EXP Bit Map Forwarding Class    Profile
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Dot1p Bit Map                 Forwarding Class              Profile
-------------------------------------------------------------------------------
3                                      ef                   n
4                                      af                   Out
5                                      nc                   Use-DE
-------------------------------------------------------------------------------
Egress Forwarding Class Queuing
-------------------------------------------------------------------------------
FC Value       : 0                              FC Name      : be
- DSCP Mapping
Out-of-Profile : be                             In-Profile   : be

- Dot1p Mapping
Out-of-Profile : 7                              In-Profile   : 7

- LSP EXP Bit Mapping
Out-of-Profile : 0                              In-Profile   : 0

- DE Mark      : Force 1

FC Value       : 1                              FC Name      : l2
- DSCP Mapping
Out-of-Profile : cs1                            In-Profile   : cs1
```

```
- Dot1p Mapping
Out-of-Profile : 1                              In-Profile   : 1

- LSP EXP Bit Mapping
Out-of-Profile : 1                              In-Profile   : 1

- DE Mark      : None
-------------------------------------------------------------------------
config>qos#
```

sgt-qos

Syntax **sgt-qos**

Context show>router

Description This command displays self-generated traffic QoS related information.

application

Syntax **application** [*app-name*] [**dscp|dot1p**]

Context show>router>sgt-qos

Description This command displays application QoS settings.

Parameters *app-name* — The specific application.

Values arp, bgp, cflowd, dhcp, dns, ftp, icmp, igmp, isis, ldp, mld, msdp, ndis, ntp, ospf, pimradius, rip, rsvpsnmp, snmp-notification, srrp, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp, pppoe

dscp-map

Syntax **dscp-map** [*dscp-name*]

Context show>router>sgt-qos

Description This command displays DSCP to FC mappings.

Parameters *dscp-name* — The specific DSCP name.

be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

Network Queue QoS Policies

In This Section

This section provides information to configure network queue QoS policies using the command line interface.

Topics in this section include:

- Overview on page 146
- Basic Configurations on page 147
- Default Network Queue Policy Values on page 153
- Service Management Tasks on page 159

Overview

Network queue policies define the ingress network queuing at the MDA network node level. Network queue policies are also used at the Ethernet port and SONET/SDH path level to define network egress queuing.

There is one default network queue policy. Each policy can have up to 16 queues (unicast and multicast). The default policies can be copied but they cannot be deleted or modified. The default policy is identified as **network-queue default**. Default network queue policies are applied to MDA network ingress ports. You must explicitly create and then associate other network queue QoS policies.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your devices, refer to CLI Usage chapter in the Basic System Configuration Guide.

Network Queue Parent Scheduler

Network queues support port scheduler parent priority-level associations. Using a port scheduler policy definition and mapping network queues to a port parent priority level, HQoS functionality is supported providing eight levels of strict priority and weights within the same priority. A network queue's bandwidth is allocated using the "within-cir" and "above-cir" scheme normal for port schedulers.

Queue CIR and PIR percentages when port-based schedulers are in effect will be based on frame-offered-load calculations.

A network queue with a port parent association exists on a port without a scheduler policy defined will be considered to be orphaned.

Refer to QoS Scheduler Policies on page 499 for more information about queue parental association scope.

Basic Configurations

A basic network queue QoS policy must conform to the following:

- Each network queue QoS policy must have a unique policy name.
- Queue parameters can be modified, but cannot be deleted.

Create a Network Queue QoS Policy

Configuring and applying QoS policies other than the default policy is optional. A default network queue policy is applied to MDA network ingress ports.

To create an network queue policy, define the following:

- Enter a network queue policy name. The system will not dynamically assign a name.
- Include a description. The description provides a brief overview of policy features.
- Forwarding class — You can assign a forwarding class to a specific queue.

Use the following CLI syntax to create a network queue QoS policy:

CLI Syntax:

```
config>qos
    network-queue policy-name
        description description-string
        fc fc-name
            multicast-queue queue-id
            queue queue-id
        queue queue-id [multipoint] [queue-type]
            cbs percent
            high-prio-only percent
            mbs percent
            port-parent [weight weight] [level level] [cir-weight
                cir-weight] [cir-level cir-level]
            rate percent [cir percent]
```

```
A:ALA-1>config>qos# network-queue default
A:ALA-1>config>qos>network-queue# info detail
----------------------------------------------
            description "Default network queue QoS policy."
            queue 1 create
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 2 create
                rate 100 cir 25
                mbs 50
```

```
        cbs 3
        high-prio-only 10
    exit
    queue 3 create
        rate 100 cir 25
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 4 create
        rate 100 cir 25
        mbs 25
        cbs 3
        high-prio-only 10
    exit
    queue 5 create
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 6 create
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 7 create
        rate 100 cir 10
        mbs 25
        cbs 3
        high-prio-only 10
    exit
    queue 8 create
        rate 100 cir 10
        mbs 25
        cbs 3
        high-prio-only 10
    exit
    queue 9 multipoint create
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 10 multipoint create
        rate 100 cir 5
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 11 multipoint create
        rate 100 cir 5
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 12 multipoint create
        rate 100 cir 5
        mbs 25
```

```
        cbs 1
        high-prio-only 10
    exit
    queue 13 multipoint create
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 14 multipoint create
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 15 multipoint create
        rate 100 cir 10
        mbs 25
        cbs 1
        high-prio-only 10
    exit
    queue 16 multipoint create
        rate 100 cir 10
        mbs 25
        cbs 1
        high-prio-only 10
    exit
    fc af create
        multicast-queue 11
        queue 3
    exit
    fc be create
        multicast-queue 9
        queue 1
    exit
    fc ef create
        multicast-queue 14
        queue 6
    exit
    fc h1 create
        multicast-queue 15
        queue 7
    exit
    fc h2 create
        multicast-queue 13
        queue 5
    exit
    fc l1 create
        multicast-queue 12
        queue 4
    exit
    fc l2 create
        multicast-queue 10
        queue 2
    exit
    fc nc create
        multicast-queue 16
        queue 8
    exit
```

Applying Network Queue Policies

Apply network queue policies to the following entities:

- MDAs
- Ethernet Ports
- SONET/SDH Ports

MDAs

Use the following CLI syntax to apply a network queue policy to an MDA network ingress port:

CLI Syntax:
```
config>card
    mda mda-slot
        network
            ingress
                queue-policy name
```

The following output displays MDA network ingress queue policy reverted to the default policy.

```
A:ALA-7>config>card>mda# info
----------------------------------------------
            mda-type m60-10/100eth-tx
            network
                ingress
                    pool default
                        resv-cbs sum
                        slope-policy "default"
                    exit
                    queue-policy "default"
                exit
                egress
                    pool default
                        resv-cbs sum
                        slope-policy "default"
                    exit
                exit
            exit
            access
                ingress
                    pool default
                        resv-cbs sum
                        slope-policy "default"
                    exit
                exit
                egress
                    pool default
                        resv-cbs sum
```

```
                        slope-policy "default"
                    exit
                exit
            exit
            no shutdown
----------------------------------------------
A:ALA-7>config>card>mda#
```

Ethernet Ports

Use the following CLI syntax to apply a network queue policy to an Ethernet port.

CLI Syntax: config>port#
ethernet
network
queue-policy *name*

```
A:ALA-49>config>port# info
----------------------------------------------
        ethernet
            network
                queue-policy "nq1"
            exit
        exit
        no shutdown
----------------------------------------------
A:ALA-49>config>port#
```

SONET/SDH Ports

Use the following CLI syntax to apply a network queue policy to a SONET/SDH port:

CLI Syntax:
```
config>port#
 sonet-sdh
    path path
       network
          queue-policy name
```

The following output displays the port configuration.

```
A:ALA-48>config>port# info
----------------------------------------------
        description "OC-12 SONET/SDH"
        sonet-sdh
            path sts3
                network
                    queue-policy "nq1"
                exit
                no shutdown
            exit
        exit
        no shutdown
----------------------------------------------
A:ALA-48>config>port#
```

Default Network Queue Policy Values

The default network queue policies are identified as policy-id **default**. The default policies cannot be modified or deleted. The following displays default policy parameters:

Table 28: Network Queue Policy Defaults

Field	Default
description	"Default network queue QoS policy."
queue 1	
pir	100
cir	0
mbs	50
cbs	1
high-prio-only	10
queue 2	
pir	100
cir	25
mbs	50
cbs	3
high-prio-only	10
queue 3	
pir	100
cir	25
mbs	50
cbs	1
high-prio-only	10
queue 4	
pir	100
cir	25

Table 28: Network Queue Policy Defaults (Continued)

Field	Default
mbs	25
cbs	3
high-prio-only	10
queue 5	
pir	100
cir	100
mbs	50
cbs	1
high-prio-only	10
queue 6	
pir	100
cir	100
mbs	50
cbs	1
high-prio-only	10
queue 7	
pir	100
cir	10
mbs	25
cbs	3
high-prio-only	10
queue 8	
pir	100
cir	10
mbs	50
cbs	3

Table 28: Network Queue Policy Defaults (Continued)

Field	Default
high-prio-only	10
fc af	queue 3
	multicast-queue 11
fc be	queue 1
	multicast-queue 9
fc ef	queue 6
	multicast-queue 14
fc h1	queue 67
	multicast-queue 15
fc h2	queue 5
	multicast-queue 13
fc l1	queue 7
	multicast-queue 12
fc l2	queue 2
	multicast-queue 10
fc nc	queue 8
	multicast-queue 16

```
A:ALA-7>config>qos>network-queue# info detail
----------------------------------------------
            description "Default network queue QoS policy."
            queue 1 auto-expedite create
                rate 100 cir 0
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 2 auto-expedite create
                rate 100 cir 25
                mbs 50
                cbs 3
                high-prio-only 10
            exit
            queue 3 auto-expedite create
                rate 100 cir 25
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 4 auto-expedite create
                rate 100 cir 25
                mbs 25
                cbs 3
                high-prio-only 10
            exit
            queue 5 auto-expedite create
                rate 100 cir 100
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 6 auto-expedite create
                rate 100 cir 100
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 7 auto-expedite create
                rate 100 cir 10
                mbs 25
                cbs 3
                high-prio-only 10
            exit
            queue 8 auto-expedite create
                rate 100 cir 10
                mbs 25
                cbs 3
                high-prio-only 10
            exit
            queue 9 multipoint auto-expedite create
                rate 100 cir 0
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 10 multipoint auto-expedite create
                rate 100 cir 5
```

```
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 11 multipoint auto-expedite create
        rate 100 cir 5
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 12 multipoint auto-expedite create
        rate 100 cir 5
        mbs 25
        cbs 1
        high-prio-only 10
    exit
    queue 13 multipoint auto-expedite create
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 14 multipoint auto-expedite create
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 15 multipoint auto-expedite create
        rate 100 cir 10
        mbs 25
        cbs 1
        high-prio-only 10
    exit
    queue 16 multipoint auto-expedite create
        rate 100 cir 10
        mbs 25
        cbs 1
        high-prio-only 10
    exit
    fc af create
        multicast-queue 11
        queue 3
    exit
    fc be create
        multicast-queue 9
        queue 1
    exit
    fc ef create
        multicast-queue 14
        queue 6
    exit
    fc h1 create
        multicast-queue 15
        queue 7
    exit
    fc h2 create
        multicast-queue 13
        queue 5
```

```
                exit
                fc l1 create
                    multicast-queue 12
                    queue 4
                exit
                fc l2 create
                    multicast-queue 10
                    queue 2
                exit
                fc nc create
                    multicast-queue 16
                    queue 8
                exit
------------------------------------------------
A:ALA-7>config>qos>network-queue#
```

Service Management Tasks

This section discusses the following service management tasks:

- Deleting QoS Policies on page 159
- Remove a Policy from the QoS Configuration on page 159
- Copying and Overwriting QoS Policies on page 160
- Editing QoS Policies on page 165

Deleting QoS Policies

A network queue policy is associated by default with MDA network ingress ports. You can replace the default policy with a customer-configured policy, but you cannot entirely remove a QoS policy. When you remove a QoS policy, the policy association reverts to the default network-queue policy **default**.

To delete a user-created network queue policy, enter the following commands:

CLI Syntax: `config>qos# no network-queue` *`policy-name`*

Example: `config>qos# no network-queue` ***`nq1`***

Remove a Policy from the QoS Configuration

To delete a network policy, enter the following commands:

CLI Syntax: `config>qos# no network-queue` *`policy-name`*

Example: `config>qos# no network-queue` *`test`*

Copying and Overwriting QoS Policies

You can copy an existing network queue policy, rename it with a new policy ID name, or overwrite an existing network queue policy. The `overwrite` option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos# copy network-queue` *`source-policy-id dest-policy-id`* `[overwrite]`

Example:`config>qos# copy network-queue` **`nq1 nq2`**

The following output displays the copied policies:

```
A:ALA-12>config>qos# info
#------------------------------------------
echo "QoS Slope/Queue Policies Configuration"
#------------------------------------------
...
        network-queue "nq1" create
            description "Default network queue QoS policy."
            queue 1 create
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 2 create
                rate 100 cir 25
                mbs 50
                cbs 3
                high-prio-only 10
            exit
            queue 3 create
                rate 100 cir 25
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 4 create
                rate 100 cir 25
                mbs 25
                cbs 3
                high-prio-only 10
            exit
            queue 5 create
                rate 100 cir 100
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 6 create
                rate 100 cir 100
                mbs 50
                cbs 1
                high-prio-only 10
            exit
```

```
queue 7 create
    rate 100 cir 10
    mbs 25
    cbs 3
    high-prio-only 10
exit
queue 8 create
    rate 100 cir 10
    mbs 25
    cbs 3
    high-prio-only 10
exit
queue 9 multipoint create
    mbs 50
    cbs 1
    high-prio-only 10
exit
queue 10 multipoint create
    rate 100 cir 5
    mbs 50
    cbs 1
    high-prio-only 10
exit
queue 11 multipoint create
    rate 100 cir 5
    mbs 50
    cbs 1
    high-prio-only 10
exit
queue 12 multipoint create
    rate 100 cir 5
    mbs 25
    cbs 1
    high-prio-only 10
exit
queue 13 multipoint create
    rate 100 cir 100
    mbs 50
    cbs 1
    high-prio-only 10
exit
queue 14 multipoint create
    rate 100 cir 100
    mbs 50
    cbs 1
    high-prio-only 10
exit
queue 15 multipoint create
    rate 100 cir 10
    mbs 25
    cbs 1
    high-prio-only 10
exit
queue 16 multipoint create
    rate 100 cir 10
    mbs 25
    cbs 1
    high-prio-only 10
exit
```

```
        fc af create
            multicast-queue 11
            queue 3
        exit
        fc be create
            multicast-queue 9
            queue 1
        exit
        fc ef create
            multicast-queue 14
            queue 6
        exit
        fc h1 create
            multicast-queue 15
            queue 7
        exit
        fc h2 create
            multicast-queue 13
            queue 5
        exit
        fc l1 create
            multicast-queue 12
            queue 4
        exit
        fc l2 create
            multicast-queue 10
            queue 2
        exit
        fc nc create
            multicast-queue 16
            queue 8
        exit
    exit
    network-queue "nq2" create
        description "Default network queue QoS policy."
        queue 1 create
            mbs 50
            cbs 1
            high-prio-only 10
        exit
        queue 2 create
            rate 100 cir 25
            mbs 50
            cbs 3
            high-prio-only 10
        exit
        queue 3 create
            rate 100 cir 25
            mbs 50
            cbs 1
            high-prio-only 10
        exit
        queue 4 create
            rate 100 cir 25
            mbs 25
            cbs 3
            high-prio-only 10
        exit
        queue 5 create
```

```
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 6 create
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 7 create
        rate 100 cir 10
        mbs 25
        cbs 3
        high-prio-only 10
    exit
    queue 8 create
        rate 100 cir 10
        mbs 25
        cbs 5
        high-prio-only 10
    exit
    queue 9 multipoint create
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 10 multipoint create
        rate 100 cir 5
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 11 multipoint create
        rate 100 cir 5
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 12 multipoint create
        rate 100 cir 5
        mbs 25
        cbs 1
        high-prio-only 10
    exit
    queue 13 multipoint create
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 14 multipoint create
        rate 100 cir 100
        mbs 50
        cbs 1
        high-prio-only 10
    exit
    queue 15 multipoint create
```

```
                    rate 100 cir 10
                    mbs 25
                    cbs 1
                    high-prio-only 10
                exit
                queue 16 multipoint create
                    rate 100 cir 10
                    mbs 25
                    cbs 1
                    high-prio-only 10
                exit
                fc af create
                    multicast-queue 11
                    queue 3
                exit
                fc be create
                    multicast-queue 9
                    queue 1
                exit
                fc ef create
                    multicast-queue 14
                    queue 6
                exit
                fc h1 create
                    multicast-queue 15
                    queue 7
                exit
                fc h2 create
                    multicast-queue 13
                    queue 5
                exit
                fc l1 create
                    multicast-queue 12
                    queue 4
                exit
                fc l2 create
                    multicast-queue 10
                    queue 2
                exit
                fc nc create
                    multicast-queue 16
                    queue 8
                exit
            exit
...
-------------------------------------------
A:ALA-12>config>qos#
```

Editing QoS Policies

You can change existing policies, except the default policies, and entries in the CLI. The changes are applied immediately to all interfaces where the policy is applied. To prevent configuration errors use the copy command to make a duplicate of the original policy to a work area, make the edits, and then overwrite the original policy.

Network Queue QoS Policy Command Reference

Command Hierarchies

- Configuration Commands on page 141
- Operational Commands on page 142
- Show Commands on page 142

Configuration Commands

```
config
    — qos
        — network-queue policy-name
            — description description-string
            — no description
            — [no] fc fc-name
                — multicast-queue queue-id
                — no multicast-queue
                — queue queue-id
                — no queue
                    — [no] queue
            — queue queue-id [multipoint] [queue-type] [create]
            — no queue queue-id
                — adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
                — no adaptation-rule
                — avg-frame-overhead percent
                — no avg-frame-overhead
                — cbs percent
                — no cbs
                — high-prio-only percent
                — no high-prio-only
                — mbs percent
                — no mbs
                — port-parent [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level] [cir-level level] [cir-weight weight]
                — no port-parent
                — rate percent [cir percent]
                — no rate
```

Operational Commands

```
config
    — qos
        — copy network-queue src-name dst-name [overwrite]
```

Show Commands

```
show
```

— **qos**
— **network-queue** [*network-queue-policy-name*] [**detail**]

Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
no description

Context config>qos>shared-queue
config>qos>network-queue
config>qos>network
config>qos>sap-egress
config>qos>sap-ingress
config>qos>sap-ingress>ipv6-criteria>entry
config>qos>sap-ingress>ip-criteria>entry
config>qos>sap-ingress>mac-criteria>entry
config>qos>scheduler-policy
config>qos>scheduler-policy>tier>scheduler

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax **copy network-queue** *src-name dst-name* [**overwrite**]

Context config>qos

Description This command copies or overwrites existing network queue QoS policies to another network queue policy ID.

The **copy** command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters **network-queue —** Indicates that the source policy ID and the destination policy ID are network-queue policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

overwrite — specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, a message is generated saying that the destination policy ID exists.

```
SR12>config>qos# copy network-queue nq1 nq2
MINOR: CLI Destination "nq2" exists - use {overwrite}.
SR12>config>qos# copy network-queue nq1 nq2 overwrite
```

Network Queue QoS Policy Commands

network-queue

Syntax [**no**] **network-queue** *policy-name*

Context config>qos

Description This command creates a context to configure a network queue policy. Network queue policies define the ingress network queuing at the MDA network node level and on the Ethernet port and SONET/SDH path level to define network egress queuing.

Default default

Parameters *policy-name* — The name of the network queue policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

fc

Syntax [**no**] **fc** *fc-name*

Context config>qos>network-queue

Description The **fc** context in the network-queue context provides a forwarding class queue context to the contained buffer control and queue rate commands.

The **fc** node contains the PIR, CIR, CBS and MBS commands used to control the buffer pool resources of each forwarding class queue on the ingress and egress pools that are associated with the network-queue policy.

The **no** form of this command restores all PIR, CIR, CBS and MBS parameters for the forwarding class network queue to their default values.

Parameters *fc-name* — The forwarding class name for which the contained PIR, CIR, CBS and MBS queue attributes apply. An instance of **fc** is allowed for each **fc-name**.

Values be, l2, af, l1, h2, ef, h1, nc

multicast-queue

Syntax **multicast-queue** *queue-id*
no multicast-queue

Context config>qos>network-queue>fc

Description This command overrides the default multicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all multicast traffic using this policy is forwarded using the *queue-id*.

The multicast forwarding type includes the **unknown** unicast forwarding type and the **broadcast** forwarding type unless each is explicitly defined to a different multipoint queue. When the unknown and broadcast forwarding types are left as default, they will track the defined queue for the multicast forwarding type.

The **no** form of the command sets the multicast forwarding type *queue-id* back to the default queue for the forwarding class. If the **broadcast** and **unknown** forwarding types were not explicitly defined to a multipoint queue, they will also be set back to the default multipoint queue (queue 11).

Resource Utilization:
When a multipoint queue is created and at least one forwarding class is mapped to the queue using the **multipoint-queue** command, a single ingress multipoint hardware queue is created per instance of the applied network-queue policy using the queue-policy command at the ingress network MDA level. Multipoint queues are not created at egress and the multipoint queues defined in the network-queue policy are ignored when the policy is applied to an egress port.

Parameters *queue-id* — The *queue-id* parameter specified must be an existing, multipoint queue defined in the **config>qos>network-queue>queue** context.

Values Any valid multipoint queue-ID in the policy including 2 through 16.

Default 11

queue

Syntax [**no**] **queue** *queue-id*

Context config>qos>network-queue>fc
config>qos>network-queue
config>qos>network-queue>hsmda

Description This command creates the context to configure forwarding-class to queue mappings.

Explicit definition of an ingress queue's hardware scheduler status is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or non-expedited nature of the queue based on the forwarding classes mapped to it. As long as all forwarding

classes mapped to the queue are expedited (nc, ef, h1 or h2), the queue is treated as an expedited queue by the hardware schedulers. When any non-expedited forwarding classes are mapped to the queue (be, af, l1 or l2), the queue is treated as best effort (be) by the hardware schedulers. The hardware status of the queue must be defined at the time of queue creation within the policy.

The **no** form of this command removes the *queue-id* from the network-queue policy and from any existing MDA or port using the policy. If any forwarding class forwarding types are mapped to the queue, they revert to their default queues. When a queue is removed, any pending accounting information for each MDA or port queue created due to the definition of the queue in the policy is discarded.

Resource Utilization
When the network-queue policy is applied on a ingress MDA, each unicast queue is created multiple times - once for each switch fabric destination currently provisioned. Some IOM types represent one switch fabric destinations while others may represent two. At egress, a single queue is created since the policy is applied at the port level. Queues are only created when at least one forwarding class is mapped to the queue using the queue command within the forwarding class context.

Parameters *queue-id* — The *queue-id* for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

Values 1 — 32

queue-type — The **expedite**, **best-effort** and **auto-expedite** queue types are mutually exclusive to each other. Each defines the method that the system uses to service the queue from a hardware perspective. While parental virtual schedulers can be defined for the queue, they only enforce how the queue interacts for bandwidth with other queues associated with the same scheduler hierarchy. An internal mechanism that provides access rules when the queue is vying for bandwidth with queues in other virtual schedulers is also needed. A keyword must be specified at the time the queue is created in the network-queue policy. If an attempt to change the keyword after the queue is initially defined, an error is generated.

expedite — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

auto-expedite — This keyword allows the system to auto-define the way the queue is serviced by the hardware. When **auto-expedite** is defined on the queue, the queue is treated in an expedited manner when all forwarding classes mapped to the queue are configured as expedited types nc, ef, h1 or h2. When a single non-expedited forwarding class is mapped to the queue (be, af, l1 and l2) the queue automatically falls back to non-expedited status.

Values expedite, best-effort, auto-expedite

Default auto-expedite

multipoint — This keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast

ingress traffic. If you attempt to map forwarding class unicast traffic to a multipoint queue, an error is generated and no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a pre-existing multipoint queue to edit *queue-id* parameters.

Values multipoint or not present

Default Present (the queue is created as non-multipoint)

Network Queue QoS Policy Queue Commands

queue

Syntax **queue** *queue-id* [**multipoint**] [*queue-type*] [*queue-mode*] **pool** *pool-name*
no queue *queue-id*

Context config>qos>network-queue

Description This command enables the context to configure a QoS network-queue policy queue.

Explicit definition of an ingress queue's hardware scheduler status is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or non-expedited nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (`nc`, `ef`, `h1` or `h2`), the queue is treated as an expedited queue by the hardware schedulers. When any non-expedited forwarding classes are mapped to the queue (`be`, `af`, `l1` or `l2`), the queue is treated as best effort (`be`) by the hardware schedulers. The expedited hardware schedulers are used to enforce expedited access to internal switch fabric destinations. The hardware status of the queue must be defined at the time of queue creation within the policy.

The **queue** command allows the creation of multipoint queues. Only multipoint queues can receive ingress packets that need flooding to multiple destinations. By separating the unicast for multipoint traffic at service ingress and handling the traffic on separate multipoint queues, special handling of the multipoint traffic is possible. Each queue acts as an accounting and (optionally) shaping device offering precise control over potentially expensive multicast, broadcast and unknown unicast traffic. Only the back-end support of multipoint traffic (between the forwarding class and the queue based on forwarding type) needs to be defined. The individual classification rules used to place traffic into forwarding classes are not affected. Queues must be defined as multipoint at the time of creation within the policy.

The multipoint queues are for multipoint traffic.

The **no** form of this command removes the *queue-id* from the network-queue policy and from any existing SAPs using the policy. If any forwarding class forwarding types are mapped to the queue, they revert to their default queues. When a queue is removed, any pending accounting information for each SAP queue created due to the definition of the queue in the policy is discarded.

Parameters *queue-id* — The *queue-id* for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

Values 1 — 16

queue-type — The **expedite**, **best-effort** and **auto-expedite** queue types are mutually exclusive to each other. Each defines the method that the system uses to service the queue from a hardware perspective. While parental virtual schedulers can be defined for the queue, they only enforce how the queue

interacts for bandwidth with other queues associated with the same scheduler hierarchy. An internal mechanism that provides access rules when the queue is vying for bandwidth with queues in other virtual schedulers is also needed. A keyword must be specified at the time the queue is created in the network-queue policy. If an attempt is made to change the keyword after the queue is initially defined, an error is generated.

expedite — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

auto-expedite — This keyword allows the system to auto-define the way the queue is serviced by the hardware. When **auto-expedite** is defined on the queue, the queue is treated in an expedited manner when all forwarding classes mapped to the queue are configured as expedited types `nc`, `ef`, `h1` or `h2`. When a single non-expedited forwarding class is mapped to the queue (`be`, `af`, `l1` and `l2`) the queue automatically falls back to non-expedited status.

Values expedite, best-effort, auto-expedite

Default auto-expedite

multipoint — This keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. If you attempt to map forwarding class unicast traffic to a multipoint queue, an error is generated and no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a pre-existing multipoint queue to edit *queue-id* parameters.

Values multipoint or not present

Default Not present (the queue is created as non-multipoint)

queue-mode — Specifies the mode in which the queue is operating. This attribute is associated with the queue at the time of creation and cannot be modified thereafter.

Values **profile-mode**: When the queue is operating in the profile mode (or, the color aware mode), the queue tries to provide the appropriate bandwidth to the packets with different profiles. The profiles are assigned according to the configuration of the forwarding class or the sub-forwarding class.

priority-mode: The queue is capable of handling traffic differently with two distinct priorities. These priorities are assigned by the stages preceding the queueing framework in the system. In priority mode, the queue does not have the functionality to support the profiled traffic and in such cases the queue will have a degraded performance. However, the converse is not valid and a queue in profile mode should be capable of supporting the different priorities of traffic.

Default **priority-mode**

adaptation-rule

Syntax **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
no adaptation-rule

Context config>qos>network-queue>queue

Description This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **pir** and **cir** apply.

Default adaptation-rule pir closest cir closest

Parameters *adaptation-rule* — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.

Values **pir** — Defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **pir** command is not specified, the default applies.

cir — Defines the constraints enforced when adapting the CIR rate defined within the **queue** queue-id **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

max — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

Context config>qos>network-queue>queue

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic

on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a SONET or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-Load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet-based offered-load.
- Frame-encapsulation overhead — Using the avg-frame-overhead parameter, the frame-encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10,000 octets and the avg-frame-overhead equals 10%, the frame-encapsulation overhead would be 10,000 x 0.1 or 1,000 octets.

For egress Ethernet queues, the frame-encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame-encapsulation overhead would be 50 x 20 or 1,000 octets.

- Frame-based offered-load — The frame-based offered-load is calculated by adding the offered-load to the frame-encapsulation overhead. If the offered-load is 10,000 octets and the encapsulation overhead was 1,000 octets, the frame-based offered-load would equal 11,000 octets.
- Packet to frame factor — The packet-to-frame factor is calculated by dividing the frame-encapsulation overhead by the queue's offered-load (packet based). If the frame-encapsulation overhead is 1,000 octets and the offered-load is 10,000 octets then the packet to frame factor would be 1,000 / 10,000 or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- Frame-based CIR — The frame-based CIR is calculated by multiplying the packet to frame factor with the queue's-configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame-based CIR would be 500 x 1.1 or 550 octets.
- Frame-based within-cir offered-load — The frame-based within-cir offered-load is the portion of the frame-based offered-load considered to be within the frame-based CIR. The frame-based within-cir offered-load is the lesser of the frame-based offered-load and the frame-based CIR. If the frame-based offered-load equaled 11000 octets and the frame-based CIR equaled 550 octets, the frame-based within-cir offered-load would be limited to 550 octets. If the frame-based offered-load equaled 450 octets and the frame-based CIR equaled 550 octets, the frame-based within-cir offered-load would equal 450 octets (or the entire frame-based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame-based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- Frame-based PIR — The frame-based PIR is calculated by multiplying the packet to frame-factor with the queue's-configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame-factor equals 0.1, the frame-based PIR would be 7,500 x 1.1 or 8,250 octets.
- Frame-based within-pir offered-load — The frame-based within-pir offered-load is the portion of the frame-based offered-load considered to be within the frame-based PIR. The frame-based within-pir offered-load is the lesser of the frame-based offered-load and the frame-based PIR. If the frame-based offered-load equaled 11,000 octets and the frame-based PIR equaled 8250 octets, the frame-based within-pir offered-load would be limited to 8,250 octets. If the frame-based offered-load equaled 7,000 octets and the frame-based PIR equaled 8,250 octets, the frame-based within-pir offered load would equal 7,000 octets.

Port Scheduler Operation Using Frame Transformed Rates — The port scheduler uses the frame based rates to figure the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and Subscriber SLA-Profile Average Frame Overhead Override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress-defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default **0**

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0.00 — 100.00

cbs

Syntax **cbs** *percent*
no cbs

Context config>qos>network-queue>queue

Description The Committed Burst Size (**cbs**) command specifies the relative amount of reserved buffers for a specific ingress network MDA forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.

The CBS for a queue is used to determine whether it has exhausted its reserved buffers while enqueuing packets. Once the queue has exceeded the amount of buffers considered in reserve for this queue, it must contend with other queues for the available shared buffer space within the buffer pool. Access to this shared pool space is controlled through Random Early Detection (RED) slope application.

Two RED slopes are maintained in each buffer pool. A high priority slope is used by in-profile packets. A low priority slope is used by out-of-profile packets. All Network-Control and Management packets are considered in-profile. Assured packets are handled by their in-profile and out-of-profile markings. All Best-Effort packets are considered out-of-profile. Premium queues should be configured such that the CBS percent is sufficient to prevent shared buffering of packets. This is generally taken care of by the CIR scheduling of Premium queues and the overall small amount of traffic on the class. Premium queues in a properly designed system will drain before all others, limiting their buffer utilization.

The RED slopes will detect congestion conditions and work to discard packets and slow down random TCP session flows through the queue. The RED slope definitions can be defined, modified or disabled through the network-queue policy assigned to the MDA for the network ingress buffer pool or assigned to the network port for network egress buffer pools.

The resultant CBS size can be larger than the MBS. This will result in a portion of the CBS for the queue to be unused and should be avoided.

The **no** form of this command returns the CBS size for the queue to the default for the forwarding class.

Special Cases **Forwarding Class Queue on Egress Network Port or Channel —** For network egress, each forwarding class is supported by an egress queue on a per network port basis. These forwarding class-based queues are automatically created once a port or channel is placed in the network mode. The configuration parameters for each queue come from the applied egress network-queue policy on the network port or channel.

The **cbs** value is used to calculate the queue's CBS size based on the total amount of buffer space allocated for the buffer pool on the egress network port or channel. This buffer pool size will dynamically fluctuate based on the port or channel's egress pool size setting.

The total reserved buffers based on the total percentages can exceed 100 percent. This might not be desirable and should be avoided as a rule of thumb. If the total percentage equals or exceeds 100 percent of the buffer pool size, no buffers will be available in the shared portion of the pool. Any queue exceeding its CBS size will experience a hard drop on all packets until it drains below this threshold.

Forwarding Class Queue on Ingress MDA — For network ingress, each forwarding class is supported by an ingress queue per MDA. These forwarding class queues are automatically created once a single port or channel is placed in the network mode on the MDA and are removed once all network ports or channels are removed from the MDA (defined as access). The configuration parameters for each queue come from the applied ingress policy under the network context of the MDA.

The **cbs** value is used to calculate the queue's CBS size based on the total amount buffer space allocated for the network ingress buffer pool on the MDA. This buffer pool will dynamically fluctuate based on the sum of all ingress pool sizes for all network ports and channels on the MDA.

The total reserved buffers based on the total percentages can exceed 100 percent. This might not be desirable and should be avoided as a rule of thumb. If the total percentage equals or exceeds 100 percent of the buffer pool size, no buffers will be available in the shared portion of the pool. Any queue exceeding its CBS size will experience a hard drop on all packets until it drains below this threshold.

Default The **cbs** forwarding class defaults are listed in the table below:

Table 29: cbs forwarding class defaults

Forwarding Class	Fowarding Class Label	Default CBS
Network-Control	nc	3
High-1	h1	3
Expedited	ef	1
High-2	h2	1
Low-1	l1	3
Assured	af	1
Low-2	l2	3
Best-Effort	be	1

Parameters *percent* — The percent of buffers reserved from the total buffer pool space, expressed as a decimal integer. If 10 MB is the total buffers in the buffer pool, a value of 10 would reserve 1MB (10%) of buffer space for

the forwarding class queue. The value 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can be applied for scheduling purposes).

Values 0 — 100

high-prio-only

Syntax **high-prio-only** *percent*
no high-prio-only

Context config>qos>network-queue>queue

Description The **high-prio-only** command allows the reservation of queue buffers for use exclusively by high priority packets as a default condition for access buffer queues for this network queue policy.

The difference between the MBS size for the queue and the high priority reserve defines the threshold where low priority traffic will be discarded. The result is used on the queue to define a threshold where low priority packets are discarded, leaving the rest of the default MBS size for high priority packets only. If the current MBS for the queue is 10MBytes, a value of 5 will result in a high priority reserve on the queue of 500KBytes. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.

Modifying the current MBS for the queue through the **mbs** command will cause the default **high-prio-only** function to be recalculated and applied to the queue. The **high-prio-only** command as defined for the specific queue can be used to override the default **high-prio-only** setting as defined in the network queue policy. This prevents the **high-prio-only** command for the network queue policy from having an affect on the queue.

The **no** form of this command restores the default value.

Default The **high-prio-only** forwarding class defaults are listed in the table below.

Table 30: High-prio-only forwarding class defaults

Forwarding Class	Fowarding Class Label	Default high-prio-only
Network-Control	nc	10
High-1	h1	10
Expedited	ef	10
High-2	h2	10
Low-1	l1	10
Assured	af	10

Table 30: High-prio-only forwarding class defaults

Forwarding Class	Fowarding Class Label	Default high-prio-only
Low-2	l2	10
Best-Effort	be	10

Parameters *percent* — The amount of queue buffer space, expressed as a decimal percentage of the MBS.

Values 0 — 100, default

mbs

Syntax **mbs** *percent*
no mbs

Context config>qos>network-queue>queue

Description The Maximum Burst Size (**mbs**) command specifies the relative amount of the buffer pool space for the maximum buffers for a specific ingress network MDA forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.

The MBS value is used to by a queue to determine whether it has exhausted its total allowed buffers while enqueuing packets. Once the queue has exceeded its maximum amount of buffers, all packets are discarded until the queue transmits a packet. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of the network queues.

The MBS size can sometimes be smaller than the CBS. This will result in a portion of the CBS for the queue to be unused and should be avoided.

The **no** form of the command returns the MBS size for the queue to the default for the forwarding class.

Special Cases **Forwarding Class Queue on Egress Network Port or Channel —** For network egress, each forwarding class is supported by an egress queue on a per network port basis. These forwarding class-based queues are automatically created once a port or channel is placed in the network mode. The configuration parameters for each queue come from the applied egress policy on the network port or channel.

The **mbs** value is used to calculate the queue's MBS size based on the total amount buffer space allocated for the buffer pool on the egress network port or channel. This buffer pool size will dynamically fluctuate based on the port or channels egress pool size setting.

The total MBS settings for all network egress queues on the port or channel based on the total percentages can exceed 100 percent. Some oversubscription can be desirable to allow exceptionally busy forwarding classes more access to buffer space. The proper use of CBS settings will ensure that oversubscribing MBS settings will not starve other queues of buffers when needed.

Forwarding Class Queue on Ingress MDA — For network ingress, each forwarding class is supported by an ingress queue per MDA. These forwarding class queues are automatically created once a single port or channel is placed in the network mode on the MDA and are removed once all network ports or channels are removed from the MDA (defined as access). The configuration parameters for each queue come from the applied ingress policy under the network context of the MDA.

The **mbs** value is used to calculate the queue's MBS size based on the total amount buffer space allocated for the network ingress buffer pool on the MDA. This buffer pool will dynamically fluctuate based on the sum of all ingress pool sizes for all network ports and channels on the MDA.

The total MBS settings for all network egress queues on the port or channel based on the total percentages can exceed 100 percent. Some oversubscription can be desirable to allow exceptionally busy forwarding classes more access to buffer space. The proper use of CBS settings will ensure that oversubscribing MBS settings will not starve other queues of buffers when needed.

Parameters *percent* — The percent of buffers from the total buffer pool space for the maximum amount of buffers, expressed as a decimal integer. If 10 MB is the total buffers in the buffer pool, a value of 10 would limit the maximum queue size to 1MB (10%) of buffer space for the forwarding class queue. If the total size is increased to 20MB, the existing value of 10 would automatically increase the maximum size of the queue to 2MB.

Values 0 — 100

port-parent

Syntax **port-parent** [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
no port-parent

Context config>qos>network-queue>queue

Description This command specifies whether this queue feeds off a port-level scheduler. For the network-queue policy context, only the port-parent command is supported. When a port scheduler exists on the port, network queues without a port-parent association will be treated as an orphan queue on the port scheduler and treated according to the current orphan behavior on the port scheduler. If the

port-parent command is defined for a network queue on a port without a port scheduler defined, the network queue will operate as if a parent association does not exist. Once a port scheduler policy is associated with the egress port, the port-parent command will come into effect.

When a network-queue policy is associated with an MDA or CMA for ingress queue definition, the port-parent association of the queues are ignored.

The **no** form of this command removes a port scheduler parent association for the queue or scheduler. If a port scheduler is defined on the port which the queue or scheduler instance exists, the queue or scheduler will become orphaned.

Default no port-parent

Parameters **weight *weight* —** Defines the weight the queue or scheduler will use at the above-cir port priority level (defined by the level parameter).

Values 0 — 100

Default 1

level *level* — Defines the port priority the queue or scheduler will use to receive bandwidth for its above-cir offered-load.

Values 1 — 8 (8 is the highest priority)

Default 1

cir-weight *cir-weight* — Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0, the queue or scheduler does not receive bandwidth during the port scheduler's within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 100

Default 1

cir-level *cir-level* — Defines the port priority the queue or scheduler will use to receive bandwidth for its within-cir offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port scheduler's within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 8 (8 is the highest priority)

Default 0

rate

Syntax **rate** *percent* [**cir** *percent*]
no rate

Context config>qos>network-queue>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the percentage at which the system prioritizes the queue over other queues competing for the same bandwidth. For SAP ingress, the CIR also defines the rate that packets are considered in-profile by the system. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP ingress or SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (100, 0).

Parameters **cir** *percent* **—** Defines the percentage of the guaranteed rate allowed for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **100** is assumed. Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 0 — 100

Default 100

cir *percent* **—** Defines the percentage of the maximum rate allowed for the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer.

Values 0 — 100

Default 0

Show Commands

network-queue

Syntax **network-queue** [*network-queue-policy-name*] [**detail**]

Description This command displays network queue policy information.

Context show>qos

Parameters *network-queue-policy-name* — The name of the network queue policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

detail — Includes each queue's rates and adaptation-rule and & cbs details. It also shows FC to queue mapping details.

Table 31: Network Queue Labels and Descriptions

Label	Description
Policy	The policy name that uniquely identifies the policy.
Description	A text string that helps identify the policy's context in the configuration file.
Port-Id	Displays the physical port identifier where the network queue policy is applied.
Queue	Displays the queue ID.
CIR	Displays the committed information rate.
PIR	Displays the peak information rate.
CBS	Displays the committed burst size.
MBS	Displays the maximum burst size.
HiPrio	Displays the high priority value.
FC	Displays FC to queue mapping.
UCastQ	Displays the specific unicast queue to be used for packets in the forwarding class.

```
A:ALA-12# show qos network-queue nq1
===============================================================================
```

```
QoS Network Queue Policy
===============================================================================
Network Queue Policy (nq1)
-------------------------------------------------------------------------------
Policy         : nq1
Description    : (Not Specified)
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Port-id : 1/1/1
===============================================================================
A:ALA-12>show>qos#

A:ALA-12>show>qos# network-queue nq1 detail
===============================================================================
QoS Network Queue Policy
===============================================================================
Network Queue Policy (nq1)
-------------------------------------------------------------------------------
Policy         : nq1
Description    : (Not Specified)
-------------------------------------------------------------------------------
Queue CIR       PIR       CBS     MBS     HiPrio
-------------------------------------------------------------------------------
1     0         100       1       50      10
2     25        100       5       50      10
3     25        100       20      50      10
4     25        100       5       25      10
5     100       100       20      50      10
6     100       100       20      50      10
7     10        100       5       25      10
8     10        100       5       25      10
-------------------------------------------------------------------------------
FC              UCastQ
-------------------------------------------------------------------------------
be              1
l2              2
af              3
l1              4
h2              5
ef              6
h1              7
nc              8

-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Port-id : 1/1/1
===============================================================================
A:ALA-12>show>qos#
```

Service Egress and Ingress QoS Policies

In This Section

This section provides information to configure SAP ingress and egress QoS policies using the command line interface.

Topics in this section include:

- Overview on page 190
  - → Egress SAP Forwarding Class and Forwarding Profile Overrides on page 191
  - → DEI Egress Remarking on page 192
  - → Default Service Egress and Egress Policy Values on page 199
  - → VID Filters on page 216
- Basic Configurations on page 202
- Service Management Tasks on page 223

Overview

There is one default service ingress policy and one default service egress policy. Each policy can have up to 32 ingress queues and 8 egress queues per service.

Each policy can have up to 32 ingress queues and 8 egress queues per service. The default policies can be copied and modified but they cannot be deleted. The default policies are identified as policy ID 1.

The default policies are applied to the appropriate interface, by default. For example, the default SAP ingress policy is applied to access ingress SAPs.The default SAP egress policy is applied to access egress SAPs. You must explicitly associate other QoS policies.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your router, refer to the CLI Usage chapter in the Basic System Configuration Guide.

Egress SAP Forwarding Class and Forwarding Profile Overrides

An access egress packet's forwarding class can be changed to redirect the packet to an alternate queue than the ingress forwarding class determination would have used. An access egress packet's profile (in or out) can also be changed to modifying the congestion behavior within the egress queue. In both cases, egress marking decisions will be based on the new forwarding class and profile as opposed to the egress forwarding class or profile. The exception is when ingress remarking is configured. An ingress remark decision will not be affected by egress forwarding class or egress profile overrides.

SAP Egress QoS Policy Modifications

The SAP egress QoS policy allows reclassification rules that are used to override the ingress forwarding class and profile of packets that egress a SAP where the QoS policy is applied. Only IP-based reclassification rules are supported.

IP precedence, DSCP and IP quintuple entries can be defined, each with an explicit forwarding class or profile override parameters. The reclassification logic for each entry follows the same basic hierarchical behavior as the classification rules within the SAP ingress QoS policy. IP precedence and DSCP have the lowest match priority while the IP criteria (quintuple) entries have the highest. When an optional parameter (such as **profile**) for IP precedence or DSCP entries is not specified, the value from the lower priority IP quintuple match for that parameter is preserved. If the IP precedence values overlap with DSCP values in that they will match the same IP header TOS field, the DSCP entry parameters will override or remove the IP precedence parameters. When none of the matched entries override a parameter, the ingress classification is preserved.

Hardware Support

The egress SAP forwarding class and forwarding profile override is only supported on SAPs configured on IOM2 and IOM3 modules. If a SAP egress QoS policy with forwarding class and forwarding profile overrides are applied to a SAP on an IOM other than the IOM2 and IOM3 (such as an IOM1), no error message is generated, but the forwarding class and forwarding profile override portion of the SAP egress QoS Policy is ignored and has no effect.

DEI Egress Remarking

It is often desirable to meter traffic from different users to ensure fairness or to meet bandwidth guarantees. Dropping all traffic in excess of a committed rate is likely to result in severe under-utilization of the networks, since most traffic sources are bursty in nature. It is burdensome to meter traffic at all points in the network where bandwidth contention occurs. One solution is to mark those frames in excess of the committed rate as drop eligible on admission to the network.

Previously, the discard eligibility was marked / determined using existing QoS fields: for example, the three MPLS EXP and Ethernet dot1p bits. Using certain combination(s) of these bits to indicate both forwarding class (emission priority) and discard eligibility meant decreasing the number of Forwarding Classes that can be differentiated in the network.

IEEE 802.1ad-2005 and IEEE 802.1ah standards allow drop eligibility to be conveyed separately from priority, preserving all the eight forwarding classes (emission priorities) that could be indicated using the 3 802.1p bits. Now all the previously introduced traffic types will be marked as drop eligible. Customers can continue to use the dot1p markings with the enhancement of changing the dot1p value used, in access, based on the in/out profile information.

The following commands can be used to remark the DE values at a SAP egress:

CLI Syntax:
```
sap-egress <policy-id> create
        fc <fc-name> create
               de-mark [force <de-value>]
               de-mark-inner [force <de-value>]
               de-mark-outer [force <de-value>]
        exit
    exit
```

The precedence of the above commands is summarized as, from highest to lowest precedence:

- de-mark-outer used for outer tag markings
- de-mark-inner used for inner tag markings
- existing de-mark used for marking both tags
- markings taken from packet received at ingress

The configuration of qinq-mark-top-only under the SAP egress takes precedence over the use of the de-mark-inner in the policy, i.e. the inner VLAN tag is not remarked when qinq-mark-top-only is configured (the marking used for the inner VLAN tag is based on the current default which is governed by the marking of the packet received at the ingress to the system). If qinq-mark-top-only is omitted, both the inner and outer VLAN tags are remarked.

Note that the egress remarking occurs after any egress classification.

DEI in IEEE 802.1ad

IEEE 802.1ad-2005 standard allows drop eligibility to be conveyed separately from priority in service VLAN TAGs (STAGs) so that all of the previously introduced traffic types can be marked as drop eligible. The service VLAN TAG has a new format where the priority and discard eligibility parameters are conveyed in the three bit priority code point (PCP) field and respectively in the DE bit (Figure 10).

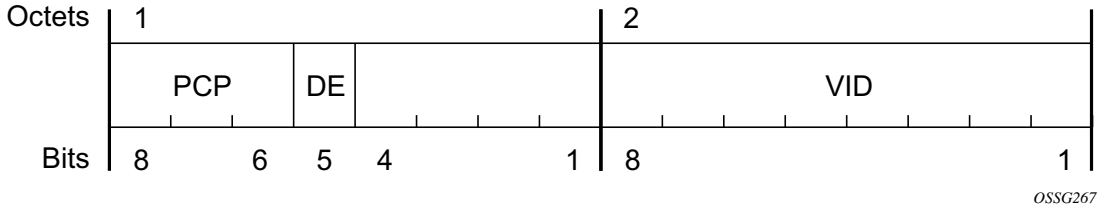

Figure 10: DE Bit in the 802.1ad S-TAG

The introduction of the DE bit allows the S-TAG to convey eight forwarding classes/distinct emission priorities, each with a drop eligible indication.

When DE bit is set to 0 (DE=FALSE) the related packet is not discard eligible. This is the case for the packets that are within the CIR limits and must be given priority in case of congestion. If the DEI is not used or backwards compliance is required the DE bit should be set to zero on transmission and ignored on reception.

When the DE bit is set to 1 (DE=TRUE) the related packet is discard eligible. This is the case for the packets that are sent above the CIR limit (but below the PIR). In case of congestion these packets will be the first ones to be dropped.

DEI in IEEE 802.1ah

IEEE 802.1ah (PBB) standard provides a dedicate bit for DE indication in both the BVID and the ITAG.

The backbone VLAN ID (BVID) is a regular 802.1ad STAG. Its DE bit may be used to convey the related tunnel QoS throughout an Ethernet backbone.

The ITAG header offers also an I-DEI bit that may be used to indicate the service drop eligibility associated with this frame.

These bits must follow the same rules as described in DEI in IEEE 802.1ad on page 193.

IEEE 802.1ad Use Case

Figure 11 illustrates an example of a topology where the new DE feature may be used: a DE aware, 802.1ad access network connected via a regular SAP to a router PE.

In this example, PE1 can ensure coherent processing of the DE indication between the 802.1ad and the MPLS networks: for example, for packets ingressing the SAP connected to 802.1ad access, read the DE indication and perform classification, color-aware metering/policing, marking of the related backbone QoS fields and selective discarding of the frames throughout the queueing system based on their discard eligibility. In addition, packets egressing the SAP towards the 802.1ad access provide proper DE indication by marking the new DE bit in the STAG.

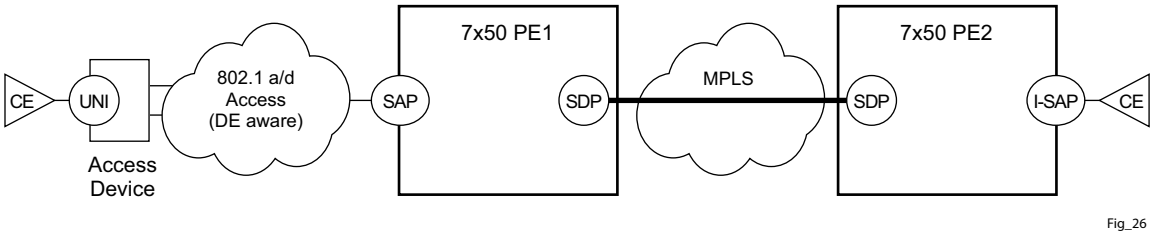

Figure 11: DE Aware 802.1ad Access Network

IEEE 802.1ah Use Case

Figure 12 illustrates an example of a PBB topology where the DE feature can be used. The processing needs highlighted in the 802.1ad use case apply to the 802.1ah BVID, format and etype being identical with the 802.1ad STAG. In addition the DE bit from the 802.1ah ITAG header may need to be processed following the same rules as for the related field in the BVID/STAG: for example, the DE bit from the BVID header represents the QoS associated with the "Ethernet Tunnel" while the DE bit from the ITAG represent the service QoS.

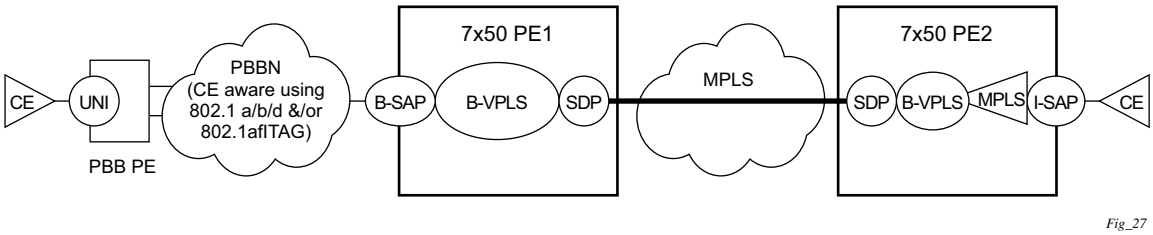

Figure 12: DE Aware PBB Topology

In this example, the BVID is not used for a part of the network leaving only I-DEI bit from the ITAG as the only option for a dedicated DE field. If both are included, then the QoS information from the BVID is to be used.

Egress FC-Based Remarking

FC-based forwarding can be used in a network using core markings of dot1p and may not support DE in all devices. The expectation is that devices beyond the network edge will continue to adhere to the end-to-end QoS policies using dot1p in the packet. Dot1p marking is performed on egress for all services and with respect to in-profile or out-of-profile context.

The following commands can be used to remark the dot1p values at a SAP egress:

CLI Syntax:

```
sap-egress <policy-id> create
        fc <fc-name> create
              dot1p {<dot1p-value>|in-profile <dot1p-value> out-
          profile <dot1p-value>}
               dot1p-inner <dot1p-value>
               dot1p-inner in-profile <dot1p-value> out-profile
          <dot1p-value>
               dot1p-outer <dot1p-value>
               dot1p-outer in-profile <dot1p-value> out-profile
```

```
            <dot1p-value>
        exit
    exit
```

The precedence of the above commands is summarized as, from highest to lowest precedence:

- dot1p-outer used for outer tag markings
- dot1p-inner used for inner tag markings
- existing dot1p used for marking both tags
- markings taken from packet received at ingress

The configuration of qinq-mark-top-only under the SAP egress takes precedence over the use of the dot1p-inner in the policy, i.e. the inner VLAN tag is not remarked when qinq-mark-top-only is configured (the marking used for the inner VLAN tag is based on the current default which is governed by the marking of the packet received at the ingress to the system). If qinq-mark-top-only is omitted, both the inner and outer VLAN tags are remarked.

Note that the egress remarking occurs after any egress classification.

Implementation Requirements

In the 7710 SR series product line, the classification to and (re-)marking from PHB (for example, forwarding class, in/out of profile status) may be described in Table 34.

Table 32: Classification to and (Re-)Marking from PHB

To/From	Classify Ingress Based on	PHB	Mark Egress To
Customer / Access Network (SAP)	dot1p [DE]	FC {in|out}	dot1p [DE]
	DSCP	FC {in|out}	DSCP
	ToS	FC {in|out}	ToS
	IP criteria	FC {in|out}	IP criteria
	MAC criteria	FC {in|out}	MAC criteria
Backbone Network (SDP / B-SAP	dot1p [DE]	FC {in|out}	dot1p [DE]
	DSCP	FC {in|out}	DSCP
	ToS	FC {in|out}	ToS
	EXP	FC {in|out}	EXP

Figure 13 displays a simple example of the DEI processing steps for the IEEE 802.1ad Use Case for both ingress and egress directions (from a PE1 SAP perspective).

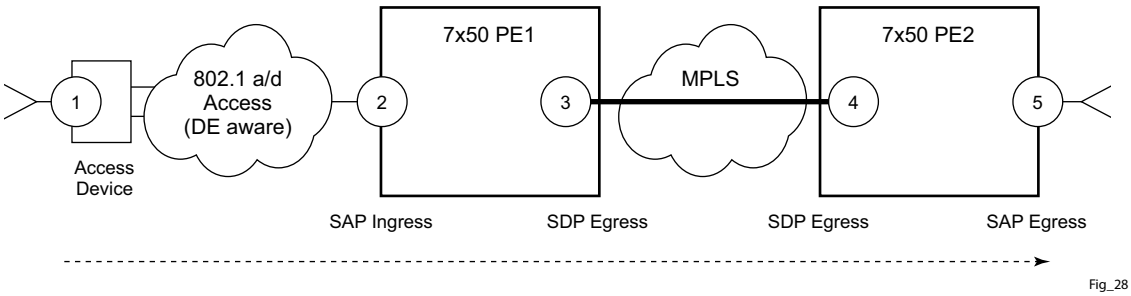

Figure 13: DEI Processing Ingress into the PE1 SAP

The following steps related to DEI are involved in the QoS processing as the packet moves from left to right:

4. The QinQ access device sets the DE bit from the STAG based on the QoS classification or on the results of the metering/policing for the corresponding customer UNI.

4. The SAP on PE1 may use the DE bit from the customer STAG to classify the frames as in/out of profile. Color aware policing/metering can generate addition out of profile packets as the result of packet flow surpassing the CIR.
5. When the packet leaves PE1 via SDP, the DE indication must be copied onto the appropriate tunnel QoS fields (outer VLAN ID and or EXP bits) using the internal PHB (per hop behavior) of the packet (for example, the FC and Profile).
6. As the packet arrives at PE2, ingress into the related SDP, the DE indication is used to classify the packets into an internal PHB.
7. Egress from the PE2 SAP, the internal PHB may be used to perform marking of the DE bit.

A combination of two access networks can be possible. If PBB encapsulation is used, the configuration used for DE in SAP and SDP policies applies to both BVID and ITAG DE bits. When both fields are used the BVID takes precedence.

Default Service Egress and Egress Policy Values

The default service egress and ingress policies are identified as policy-id **1**. The default policies cannot be edited or deleted. The following displays default policy parameters:

- SAP Egress Policy on page 199
- Default SAP Ingress Policy on page 200

SAP Egress Policy

```
A:ALA-7>config>qos>sap-egress$ info detail
----------------------------------------------
            no description
            scope template
            queue 1 auto-expedite create
                no parent
                adaptation-rule pir closest cir closest
                rate max cir 0
                cbs default
                mbs default
                high-prio-only default
            exit
----------------------------------------------
A:ALA-7>config>qos>sap-egress$
```

Table 33: SAP Egress Policy Defaults

Field	Default
description	“Default SAP egress QoS policy.”
scope	template
queue 1	1 auto-expedite
parent	no parent
adaptation-rule	adaptation-rule pir closest cir closest
rate	max cir 0
cbs	default
mbs	default
high-prio-only	default

Default SAP Ingress Policy

```
A:ALA-7>config>qos>sap-ingress$ info detail
----------------------------------------------
            description "Default SAP ingress QoS policy"
            scope template
            queue 1 auto-expedite create
                no parent
                adaptation-rule pir closest cir closest
                rate max cir 0
                mbs default
                cbs default
                high-prio-only default
            exit
            queue 2 multipoint auto-expedite create
                no parent
                adaptation-rule pir closest cir closest
                rate max cir 0
                mbs default
                cbs default
                high-prio-only default
            exit
            default-fc be
            default-priority low
----------------------------------------------
A:ALA-7>config>qos>sap-ingress$
```

Table 34: SAP Ingress Policy Defaults

Field	Default
description	"Default SAP ingress QoS policy."
scope	template
queue 1	1 priority-mode auto-expedite
parent	no parent
adaptation-rule	adaptation-rule pir closest cir closest
rate	max cir 0
cbs	default
mbs	default
high-prio-only	default
queue 2	multipoint priority-mode auto-expedite
parent	no parent
adaptation-rule	adaptation-rule pir closest cir closest
rate	max cir 0

Table 34: SAP Ingress Policy Defaults (Continued)

Field	Default
cbs	default
mbs	default
high-prio-only	default
default-fc	be
default-priority	low

Basic Configurations

A basic service egress QoS policy must conform to the following:

- Have a unique service egress QoS policy ID.
- Have a QoS policy scope of template or exclusive.
- Have at least one defined default queue.

A basic service ingress QoS policy must conform to the following:

- Have a unique service ingress QoS policy ID.
- Have a QoS policy scope of template or exclusive.
- Have at least one default unicast forwarding class queue.
- Have at least one multipoint forwarding class queue.

Create Service Egress and Ingress QoS Policies

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP or IP interface, a default QoS policy is applied.

- Percent-Rate Support on page 202
- Service Egress QoS Policy on page 204
- Service Ingress QoS Policy on page 206

Percent-Rate Support

The **percent-rate** command is supported for pir and cir parameters for both queues and policers. Also supported is the capability of specifying the rate as a percentage value of the line rate for sap-ingress and sap-egress qos policies. It is supported for both queues and policers. The user has the option of specifying **percent-rate** for **pir** and **cir** parameters. For **pir**, the range is 0.01 to 100.00, and for **cir**, the range is 0.00 to 100.00.

The rate can be also configured using the existing keyword **rate** in Kbps.

For queues, when the queue rate is in percent-rate either a local-limit or a port-limit can be applied.

When the local-limit is used the percent-rate is relative to the queue's parent scheduler rate or the agg-rate rate at egress, when the port-limit is used the percent-rate is relative to the rate of the port (including the ingress-rate/egress-rate setting) to which the queue is attached. port-limit is the default.

For policers, the percent-rate rate is always relative to the immediate parent root policer/arbiter rate or the FP capacity.

SAP Ingress QoS Policy:

```
*B:Dut-A>config>qos>sap-ingress# queue 1 percent-rate
- no percent-rate
- percent-rate <pir-percent> [cir <cir-percent>] [port-limit|local-limit]
- percent-rate <pir-percent> police [port-limit|local-limit]

<pir-percent> : [0.01..100.00]
<cir-percent> : [0.00..100.00]
<police> : keyword
<port-limit|local-*> : keyword

*B:Dut-A>config>qos>sap-ingress# policer 1 percent-rate
- no percent-rate
- percent-rate <pir-percent> [cir <cir-percent>]

<pir-percent> : [0.01..100.00]
<cir-percent> : [0.00..100.00]
```

SAP-Egress QoS Policy:

```
*B:Dut-A>config>qos>sap-egress# queue 1 percent-rate
- no percent-rate
- percent-rate <pir-percent> [cir <cir-percent>] [port-limit|local-limit]

<pir-percent> : [0.01..100.00]
<cir-percent> : [0.00..100.00]
<port-limit|local-*> : keyword

*B:Dut-A>config>qos>sap-egress# policer 1 percent-rate
- no percent-rate
- percent-rate <pir-percent> [cir <cir-percent>]

<pir-percent> : [0.01..100.00]
<cir-percent> : [0.00..100.00]
```

Service Egress QoS Policy

To create a service egress policy, you must define the following:

- A new policy ID value. The system will not dynamically assign a value.
- Specify the scope. A QoS policy must be defined as having either an *exclusive* scope for one-time use, or a *template* scope which enables its use with multiple SAPs.
- Include a description. The description provides a brief overview of policy features.

After the policy is created, the policy's behavior can be defined:

- Specify the forwarding class. The forwarding class name or names associated with the egress queue. The egress queue for the service traffic is selected based on the forwarding classes that are associated with the queue.
- A new queue ID value. The system will not dynamically assign a value.
- Define queue parameters. Queues support explicit and auto-expedite hardware queue scheduling, and parent virtual scheduler definition.

The following displays an egress QoS policy configuration:

```
A:ALA-7>config>qos# info
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
     sap-egress 105 create
            description "SAP egress policy"
            queue 1 create
            exit
            queue 2 create
            exit
            queue 3 expedite create
                parent test1
            exit
            fc af create
                queue 1
            exit
            fc ef create
                queue 2
            exit
        exit
...
#------------------------------------------
A:ALA-7>config>qos#
```

Service Egress QoS Queue

To create a service egress queue parameters, define the following:

- A new queue ID value. The system will not dynamically assign a value.
- Define queue parameters. Egress queues support explicit and auto-expedite hardware queue scheduling, and parent virtual scheduler definition.

The following displays an egress QoS policy configuration:

```
A:ALA-7>config>qos# info
----------------------------------------------
...
        sap-egress 105 create
            description "SAP egress policy"
            queue 1 create
                parent "scheduler-tier1"
            exit
            queue 2 create
            exit
            queue 3 expedite create
                parent "test1"
            exit
            fc af create
                queue 1
            exit
            fc ef create
            exit
        exit
...
----------------------------------------------
A:ALA-7>config>qos#
```

Service Ingress QoS Policy

To create an service ingress policy, define the following:

- A policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Specify a default forwarding class for the policy. All packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class.
- Specify a default priority for all packets received on an ingress SAP using this policy.
- Define mappings from incoming packet contents to a forwarding class, and then, separately, from the forwarding class to queue.
- Define forwarding class parameters.
  - → Modify the **multicast-queue** default value to override the default multicast forwarding type queues mapping for **fc** *fc-name*.
  - → Modify the **unknown-queue** default value to override the default unknown unicast forwarding type queues mapping for **fc** *fc-name*.
  - → Modify the **broadcast-queue** default value to override the default broadcast forwarding type queues mapping for **fc** *fc-name*.
- Configure precedence value for the forwarding class or enqueuing priority when a packet is marked with an IP precedence value.
- Specify IP, IPv6 or MAC criteria. You can define IP, IPv6 and MAC-based SAP ingress policies to select the appropriate ingress queue and corresponding forwarding class for matched traffic.
- A SAP ingress policy is created with a template scope. The scope can be modified to exclusive for a special one-time use policy. Otherwise, the **template** scope enables the policy to be applied to multiple SAPs.

The following displays an service ingress policy configuration:

```
A:ALA-7>config>qos>sap-ingress# info
----------------------------------------------
...
        sap-ingress 100 create
            description "Used on VPN sap"
...
----------------------------------------------
A:ALA-7>config>qos>sap-ingress#
```

Service Ingress QoS Queue

To create service ingress queues parameters, define the following:

- A new queue ID value — The system will not dynamically assign a value.
- Queue parameters — Ingress queues support multipoint queues, explicit and auto-expedite hardware queue scheduling, and parent virtual scheduler definition.

The following displays an ingress queue configuration:

```
A:ALA-7>config>qos# info
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
        sap-ingress 100 create
            description "Used on VPN sap"
            queue 1 create
            exit
            queue 2 multipoint create
            exit
            queue 10 create
                parent VPN_be
                rate 11000
            exit
            queue 12 create
                parent VPN_priority
                rate 11000
            exit
            queue 13 create
                parent VPN_reserved
                rate 1
            exit
            queue 15 create
                parent VPN_video
                rate 1500 cir 1500
            exit
            queue 16 create
                parent VPN_voice
                rate 2500 cir 2500
            exit
            queue 17 create
                parent VPN_nc
                rate 100 cir 36
            exit
            queue 20 multipoint create
                parent VPN_be
                rate 11000
            exit
            queue 22 multipoint create
                parent VPN_priority
                rate 11000
            exit
            queue 23 multipoint create
                parent VPN_reserved
```

```
                rate 1
            exit
            queue 25 multipoint create
                parent VPN_video
                rate 1500 cir 1500
            exit
            queue 26 multipoint create
                parent VPN_voice
                rate 2500 cir 2500
            exit
            queue 27 multipoint create
                parent VPN_nc
                rate 100 cir 36
            exit
...
#--------------------------------------------
A:ALA-7>config>qos#
```

SAP Ingress Forwarding Class (FC)

The following displays a forwarding class and precedence configurations:

```
A:ALA-7>config>qos# info
#------------------------------------------
...
            fc af create
                queue 12
                broadcast-queue 22
                multicast-queue 22
                unknown-queue 22
            exit
            fc be create
                queue 10
                broadcast-queue 20
                multicast-queue 20
                unknown-queue 20
            exit
            fc ef create
                queue 13
                broadcast-queue 23
                multicast-queue 23
                unknown-queue 23
            exit
            fc h1 create
                queue 15
                broadcast-queue 25
                multicast-queue 25
                unknown-queue 25
            exit
            fc h2 create
                queue 16
                broadcast-queue 26
                multicast-queue 26
                unknown-queue 26
            exit
            fc nc create
                queue 17
                broadcast-queue 27
                multicast-queue 27
                unknown-queue 27
            exit
            prec 0 fc be
            prec 2 fc af
            prec 3 fc ef
            prec 5 fc h1
            prec 6 fc h2
            prec 7 fc nc
...
#------------------------------------------
A:ALA-7>config>qos#
```

Service Ingress IP Match Criteria

When specifying SAP ingress match criteria, only one match criteria type can be configured in the SAP ingress QoS policy.

The following displays an ingress IP criteria configuration:

```
A:ALA-7>config>qos# info
...
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
        sap-ingress 100 create
...
            ip-criteria
                entry 10 create
                    description "Entry 10-FC-AF"
                    match protocol 6
                        src-ip 10.10.10.103/24
                    exit
                    action fc af priority high
                exit
                entry 20 create
                    description "Entry 20-FC-BE"
                    match protocol 17
                        dst-port eq 255
                    exit
                    no action
                exit
            exit
        exit
..
#------------------------------------------
A:ALA-7>config>qos#
```

Service Ingress IPv6 Match Criteria

When specifying SAP ingress match criteria, only one match criteria type can be configured in the SAP ingress QoS policy.

The following displays an ingress IPv6 criteria configuration:

```
A:ALA-48>config>qos>sap-ingress# info
----------------------------------------------
            queue 1 create
            exit
            queue 11 multipoint create
            exit
            ip-criteria
            exit
            ipv6-criteria
                entry 10 create
                    description "IPv6 SAP-ingress policy"
                    match
                        src-ip ::/96
                        dst-ip 200::/7
                    exit
                    action fc be priority low
                exit
                entry 20 create
                    description "Entry 20-FC-AF"
                    match next-header tcp
                        src-port eq 500
                    exit
                    action fc af priority high
                exit
            exit
----------------------------------------------
A:ALA-48>config>qos>sap-ingress#
```

Service Ingress MAC Match Criteria

Both IP criteria and MAC criteria cannot be configured in the same SAP ingress QoS policy.

To configure service ingress policy MAC criteria, define the following:

- A new entry ID value. Entries must be explicitly created. The system will not dynamically assign entries or a value.
- The action to associate the forwarding class or enqueuing priority with a specific MAC criteria entry ID.
- A description. The description provides a brief overview of policy features.
- Match criteria for ingress SAP QoS policy. Optionally, specify an IP protocol to be used as an ingress SAP QoS policy match criterion.

The following displays an ingress MAC criteria configuration:

```
A:ALA-7>config>qos# info
...
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
        sap-ingress 101 create
...
            mac-criteria
                entry 10 create
                    description "Entry10-low prio"
                    match
                        dst-mac 04-67-ff-00-00-01 ff-ff-ff-ff-ff-ff
                        dot1p 7 7
                    exit
                    action fc be priority low
                exit
            exit
        exit
#------------------------------------------
A:ALA-7>config>qos#
```

Ingress Criteria Classification Directly to Policer

It is possible to classify traffic directly to a policer, independent of the policer/queue assigned to the traffic's forwarding class. This is supported at SAP ingress when using one of the following statements: ip-criteria, ipv6-critera or mac-criteria.

The standard mechanisms are still used to assign a forwarding class to the related traffic, and this forwarding class continues to be used for QOS processing at egress.

This is supported on all FP2 and higher based line cards. The use of explicitly configured broadcast, unknown, or multicast policers is not supported. QPPB processing takes precedence over this feature.

This could be used, for example, when it is required that ingress OAM traffic is not subject to the same QOS control as other customer traffic on a given SAP. The OAM traffic could be classified based on its source MAC address (for example, with an OUI of 00-xx-yy as configured below) and directed to policer 1 while the remainder of the customer's traffic is processed using ingress queue 1. This is shown in Figure 14.

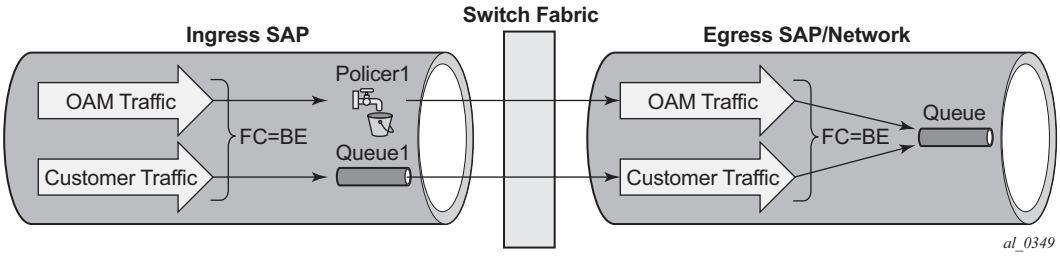

Figure 14: Ingress Criteria Classification Directly to Policer

The configuration would be as follows:

```
sap-ingress 10 create
    queue 1 create
    exit
    queue 11 multipoint create
    exit
    policer 1 create
    exit
    mac-criteria
        entry 10 create
            match
                src-mac 00-xx-yy-00-00-00 ff-ff-ff-00-00-00
            exit
            action policer 1
        exit
```

```
        exit
    exit
```

FC Mapping Based on EXP Bits

You can use the **lsp-exp** command to set your sap-ingress qos policy on Ethernet L2 SAPs to perform FC mapping based on EXP bits.

The **lsp-exp** option causes the forwarding class and drop priority of incoming traffic to be determined by the mapping result of the EXP bits in the top label.

The following example displays FC mapping based on EXP bits:

```
*A:Dut-T>config>qos>sap-ingress# info
----------------------------------------------
            queue 1 create
            exit
            queue 2 create
            exit
            queue 3 create
            exit
            queue 11 multipoint create
            exit
            fc "af" create
                queue 2
            exit
            fc "be" create
                queue 1
            exit
            fc "ef" create
                queue 3
            exit
            lsp-exp 0 fc "be" priority low
            lsp-exp 1 fc "af" priority high
            lsp-exp 2 fc "ef" priority low hsmda-counter-override 1
            lsp-exp 3 fc "ef" priority high hsmda-counter-override 2
```

VID Filters

VID filters extend the capability of current Ethernet ports with null or default SAP tag configuration to match and take action on VID tags. Service delimiting tags (for example qinq 1/1/1:10.20 or dot1q 1/1/1:10, where outer tag 10 and inner tags 20 are service delimiting) allow fine grain control of frame operations based on the VID tag. Service delimiting tags are exact match and are stripped from the frame as illustrated in Figure 15. Exact match or service delimiting tags do not require VID filters. VID filters can only be used to match on frame tags that are after the service delimiting tags.

With VID Filters operators can choose to match VID tags for up to two tags on ingress or egress or both.

- The outer-tag is the first tag in the packet that is carried transparently through the service.
- The inner-tag is the second tag in the packet that is carried transparently through the service.

VID filters add the capability to perform VID value filter policies on default tags (1/1/1:* or 1/1/1:x.*, or 1/1/1/:*.0), or null tags (1/1/1, 1/1/1:0 or 1/1/1:x.0). The matching is based on the port configuration and the SAP configuration.

QinQ tags are often referred to as the C-VID (Customer VID) and S-VID (service VID). The terms outer tag and inner tag allow flexibility without having to refer to C-TAG and an S-TAG explicitly. The position of inner and outer tags is relative to the port configuration and SAP configuration. Matching of tags is allowed for up to the first two tags on a frame. Since service delimiting tags may be 0, 1 or 2 tags.

The meaning of inner and outer has been designed to be consistent for egress and ingress when the number of non service delimiting tags is consistent. Service 1 in Figure 15 shows a conversion from qinq to a single dot1q example where there is one non-service delimiting tag on ingress and egress. Service 2 shows a symmetric example with two non-service delimiting tags (plus an additional tag for illustration) to two non-service delimiting tags on egress. Service 3 illustrates single non-service delimiting tags on ingress and to two tags with one non-service delimiting tag on ingress and egress.

SAP-ingress QoS setting allows for MAC-criteria type VID which uses the VID filter matching capabilities (see QoS and VID Filters on page 218).

A VID filter entry can be used as a debug or lawful intercept mirror source entry.

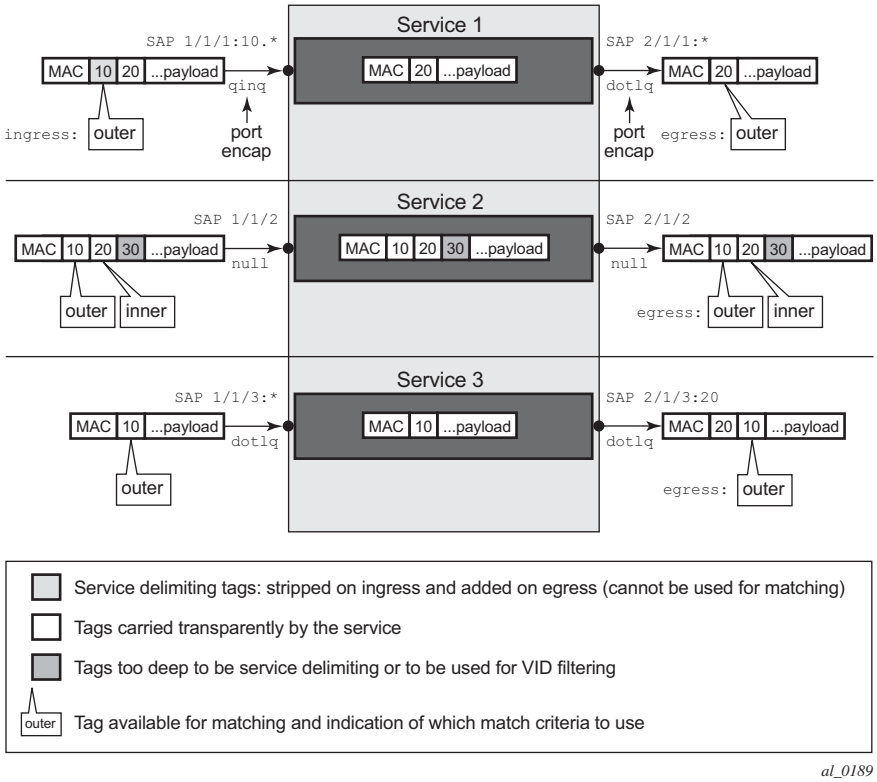

Figure 15: VID Filtering Examples

VID filters are available on Ethernet SAPs for Epipe, VPLS or I-VPLS including eth-tunnel and eth-ring services.

Arbitrary Bit Matching of VID Filters

In addition to matching an exact value, a VID filter mask allows masking any set of bits. The masking operation is ((value & vid-mask) = = (tag and vid-mask)). For example: A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6. VID filters allow explicit matching of VIDs and matching of any bit pattern within the VID tag.

When using VID filters on SAPs only VID filters are allowed on this SAP. Filters of type normal and ISID are not allowed.

An additional check for the “0” VID tag may be required when using certain wild card operations. For example frames with no tags on null encapsulated ports will match a value of 0 in outer tag and inner tag because there are no tags in the frame for matching. If a zero tag is possible but not desired it can be explicitly filtered using exact match on “0” prior to testing other bits for “0”.

Note that **configure>system>ethernet>new-qinq-untagged-sap** is a special QinQ function for single tagged QinQ frames with a null second tag. Using this in combination with VID filters is not recommended. Note that the outer-tag is the only tag available for filtering on egress for frames arriving from MPLS SDPs or from PBB services even though additional tags may be carried transparently.

QoS and VID Filters

On ingress VID filtering may also be used to set QoS on SAP ingress. The matching rules are the same as for VID filter but the action allows setting of the forwarding class.

For example, to set the forwarding class of all VIDs with 6 in the lower 3 bits of the VID a filter as illustrated below could be constructed and then ingress qos 5 could be applied to any SAP that requires the policy.

```
qos
        sap-ingress 5 create
            queue 1 create
            exit
            queue 11 multipoint create
            exit
            mac-criteria
                type vid
                entry 1 create
                    match frame-type ethernet-II
                        outer-tag 6 7
                    exit
                    action fc "af"
                exit
            exit
        exit
    exit
```

Port Group Configuration Example

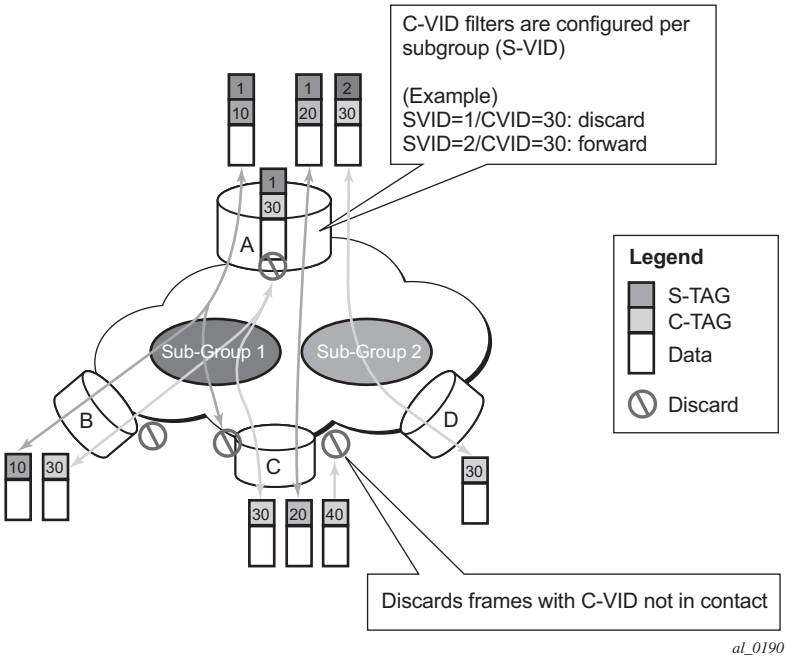

Figure 16: Port Groups

Figure 16 shows a customer use example where some VLANs are prevented from ingressing or egressing certain ports. In the example, port A sap 1/1/1:1.* would have a filter as shown below while port A sap 1/1/1:2.* would not.:

```
mac-filter 4 create
     default-action forward
            type vid
            entry 1 create
                match frame-type ethernet_II
                    outer-tag 30 4095
                exit
                action drop
            exit
        exit
```

Applying Service Ingress and Egress Policies

Apply SAP ingress and egress policies to the following service SAPs:

- Epipe
- IES
- VPLS
- VPRN

Refer to the Subscriber Services Overview section of the Services Guide for information about configuring service parameters.

Epipe

The following output displays an Epipe service configuration with SAP ingress policy 100 and SAP egress 105 applied to the SAP.

```
A:ALA-7>config>service# info
----------------------------------------------
        epipe 6 customer 6 vpn 6 create
            description "Distributed Epipe service to west coast"
            sap 1/1/10:010 create
                ingress
                    qos 100
                exit
                egress
                    qos 105
                exit
            exit
            spoke-sdp 2:6 create
                ingress
                    vc-label 6298
                exit
                egress
                    vc-label 6300
                exit
            exit
            no shutdown
        exit
----------------------------------------------
A:ALA-7>config>service#
```

IES

The following output displays an IES service configuration with SAP ingress policy 100 and SAP egress 105 applied to the SAP.

```
A:ALA-7>config>service# info
----------------------------------------------
        ies 88 customer 8 vpn 88 create
            interface "Sector A" create
                sap 1/1/1.2.2 create
                    ingress
                        qos 100
                    exit
                    egress
                        qos 105
                    exit
                exit
            exit
            no shutdown
        exit
----------------------------------------------
```

VPLS

The following output displays a VPLS service configuration with SAP ingress policy 100. The SAP egress policy 1 is applied to the SAP by default.

```
A:ALA-7>config>service# info
----------------------------------------------
        vpls 700 customer 7 vpn 700 create
            description "test"
            stp
                shutdown
            exit
            sap 1/1/9:010 create
                ingress
                    qos 100
                exit
            exit
            spoke-sdp 2:222 create
            exit
            mesh-sdp 2:700 create
            exit
            no shutdown
        exit
----------------------------------------------
A:ALA-7>config>service#
```

VPRN

The following output displays a VPRN service configuration.

```
A:ALA-7>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            ecmp 8
            autonomous-system 10000
            route-distinguisher 10001:1
            auto-bind ldp
            vrf-target target:10001:1
            interface "to-ce1" create
                address 11.1.0.1/24
                sap 1/1/10:1 create
                    ingress
                        qos 100
                    exit
                    egress
                        qos 105
                    exit
                exit
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-7>config>service#
```

Service Management Tasks

This section discusses the following service management tasks:

- Deleting QoS Policies on page 223
- Copying and Overwriting QoS Policies on page 225
- Remove a Policy from the QoS Configuration on page 226
- Editing QoS Policies on page 226

Deleting QoS Policies

Every service SAP is associated, by default, with the appropriate egress or ingress policy (policy-id **1**). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the SAP configuration. When you remove a non-default service egress or ingress policy, the association reverts to the default policy-id **1**.

A QoS policy cannot be deleted until it is removed from all SAPs where they are applied.

```
A:ALA-7>config>qos# no sap-ingress 100
MINOR: CLI SAP ingress policy "100" cannot be removed because it is in use.
A:ALA-7>config>qos#
```

Remove a QoS Policy from Service SAP(s)

The following Epipe and VPRN service output examples show that the SAP service egress and ingress reverted to policy-id "**1**" when the non-default policies were removed from the configuration.

```
A:ALA-104>config>service>epipe# info detail
----------------------------------------------
            description "Distributed Epipe service to west coast"
            service-mtu 1514
            sap 1/1/10:0 create
                no description
                no multi-service-site
                ingress
                    no scheduler-policy
                    qos 1
                exit
                egress
                    no scheduler-policy
                    qos 1
                exit
                no collect-stats
                no accounting-policy
```

```
                no shutdown
            exit
            spoke-sdp 2:6 vc-type ether create
                no shutdown
            exit
            no shutdown
-------------------------------------------
A:ALA-7>config>service>epipe#

A:ALA-7>config>service>vprn#
-------------------------------------------
...
        vprn 1 customer 1 create
            ecmp 8
            autonomous-system 10000
            route-distinguisher 10001:1
            auto-bind ldp
            vrf-target target:10001:1
            interface "to-ce1" create
                address 11.1.0.1/24
                sap 1/1/10:1 create
                exit
            exit
            no shutdown
        exit
-------------------------------------------
A:ALA-7>config>service>vprn#
```

Copying and Overwriting QoS Policies

You can copy an existing service egress or ingress policy, rename it with a new policy ID value, or overwrite an existing policy ID. The `overwrite` option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos# copy {sap-ingress | sap-egress} source-policy-id dest-policy-id [overwrite]`

The following output displays the copied policies:

```
A:ALA-7>config>qos# info
----------------------------------------------
...
exit
        sap-ingress 100 create
            description "Used on VPN sap"
            queue 1 create
            exit
            queue 2 multipoint create
            exit
            queue 10 create
                parent "VPN_be"
                rate 11000
            exit
...
        sap-ingress 101 create
            description "Used on VPN sap"
            queue 1 create
            exit
            queue 2 multipoint create
            exit
            queue 10 create
                parent "VPN_be"
                rate 11000
            exit
        sap-ingress 200 create
            description "Used on VPN sap"
            queue 1 create
            exit
            queue 2 multipoint create
            exit
            queue 10 create
                parent "VPN_be"
                rate 11000
            exit
...
----------------------------------------------
A:ALA-7>config>qos#
```

Remove a Policy from the QoS Configuration

CLI Syntax: `config>qos# no sap-ingress` *`policy-id`*

Example:
```
config>qos# no sap-ingress 100
config>qos# no sap-egress 1010
```

Editing QoS Policies

You can change QoS existing policies and entries. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors copy the policy to a work area, make the edits, and then write over the original policy.

Service SAP QoS Policy Command Reference

Command Hierarchies

- Service Ingress QoS Policy Commands
- Service Egress QoS Policy Commands
- Operational Commands
- Show Commands

Service Ingress QoS Policy Commands

config
- **qos**
  - [**no**] **sap-ingress** *policy-id* | *policy-name*
    - **default-fc** *fc-name*
    - **no default-fc**
    - **default-priority** {**low** | **high**}
    - **no default-priority**
    - **description** *description-string*
    - **no description**
    - **dot1p** *dot1p-priority* [**fc** *fc-name*] [**priority** {**low** | **high**}]
    - **no dot1p** *dot1p-priority*
    - **dscp** *dscp-name* [**fc** *fc-name*] [**priority** {**low** | **high**}]
    - **no dscp** *dscp-name*
    - [**no**] **fc** *fc-name*
      - **policer** *policer-id* [**fp-redirect-group**]
      - **no policer**
      - **broadcast-policer** *policer-id* [**fp-redirect-group**]
      - **no broadcast-policer**
      - **broadcast-queue** *queue-id* [**group** *queue-group-name*]
      - **no broadcast-queue**
      - [**no**] **de-1-out-profile**
      - **egress-fc** *fc-name*
      - **no egress-fc**
      - **in-remark dscp** *dscp-name*
      - **in-remark prec** *ip-prec-value*
      - **no in-remark**
      - **multicast-policer** *policer-id* [**fp-redirect-group**]
      - **no multicast-policer**
      - **multicast-queue** *queue-id* [**group** *queue-group-name*]
      - **no multicast-queue**
      - **out-remark dscp** *dscp-name*
      - **out-remark prec** *ip-prec-value*
      - **no out-remark**
      - **profile** {**in** | **out**}
      - **no profile**

- **queue** *queue-id* [{**group** *queue-group-name* [instance *instance-id*]} | **port-redirect-group-queue**]
- **no queue**
- **unknown-policer** *policer-id* [**fp-redirect-group**]
- **no unknown-policer**
- **unknown-queue** *queue-id* [**group** *queue-group-name*]
- **no unknown-queue**

- [**no**] **ip-criteria**
  - [**no**] **entry** *entry-id*
    - **action** [**fc** *fc-name*] [**profile** {**in** |**out**}]
    - **no action**
    - **description** *description-string*
    - **no description**
    - **match** [**protocol** *protocol-id*]
    - **no match**
      - **dscp** *dscp-name*
      - **no dscp**
      - **dst-ip** {*ip-address*/*mask* | *ip-address netmask*}
      - **no dst-ip**
      - **dst-port** {**lt** | **gt** | **eq**} *dst-port-number*
      - **dst-port range** *start end*
      - **no dst-port**
      - **fragment** {**true** | **false**}
      - **no fragment**
      - **src-ip** {*ip-address*/*prefix* | *mask* | *ip-address netmask*}
      - **no src-ip**
      - **src-port** {**lt** | **gt** | **eq**} *src-port-number*
      - **src-port range** *start end*
      - **no src-port**
  - **renum** [*old-entry-id new-entry-id*]
- [**no**] **ipv6-criteria**
  - [**no**] **entry** *entry-id*
    - **action** [**fc** *fc-name*] [**priority** {**low** | **high**}]
    - **no action**
    - **description** *description-string*
    - **no description**
    - **match** [**next-header** *next-header*]
    - **no match**
      - **dscp** *dscp-name*
      - **no dscp**
      - **dst-ip** {*ipv6-address/prefix-length*}
      - **no dst-ip**
      - **dst-port** {**lt** | **gt** | **eq**} *dst-port-number*
      - **dst-port range** *start end*
      - **no dst-port**
      - **src-ip** {*ipv6-address/prefix-length*}
      - **no src-ip**
      - **src-port** {**lt** | **gt** | **eq**} *src-port-number*
      - **src-port range** *start end*
      - **no src-port**
  - **renum** [*old-entry-id new-entry-id*]
- **lsp-exp** *lsp-exp-value* [**fc** *fc-name*] [**priority** {**low**|**high**}] [**hsmda-counter-override** *counter-id*]
- **no lsp-exp** *lsp-exp-value*

- [**no**] **mac-criteria**
  - [**no**] **entry** *entry-id*
    - **action** [**fc** *fc-name*] [**priority** {**low** | **high**}]
    - **no action**
    - **description** *description-string*
    - **no description**
    - **match** [**frame-type** {**802dot3 | 802dot2-llc | 802dot2-snap | ether-net-II | atm**}]
    - **no match**
      - **atm-vci** *vci-value*
      - **no atm-vci**
      - **dot1p** ***dot1p****-value* [*dot1p-mask*]
      - **no dot1p**
      - **dsap** *dsap-value* [*dsap-mask*]
      - **no dsap**
      - **dst-mac** *ieee-address* [*ieee-address-mask*]
      - **no dst-mac**
      - **etype** *etype-value*
      - **no etype**
      - **snap-oui** [*zero* | *non-zero*]
      - **no snap-oui**
      - **snap-pid** *snap-pid*
      - **no snap-pid**
      - **src-mac** *ieee-address* [*ieee-address-mask*]
      - **no src-mac**
      - **ssap** *ssap-value* [*ssap-mask*]
      - **no ssap**
  - **renum** *old-entry-number new-entry-number*
  -
- **policer** *policer-id* [**create**]
- **no policer** *policer-id*
  - **adaptation-rule pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
  - **no adaptation-rule**
  - **adv-config-policy** ***policy-name***
  - **no adv-config-policy**
  - **description** *description-string*
  - **no description**
  - **cbs** {*size* [**bytes** | **kilobytes**] | **default**}
  - **no cbs**
  - **high-prio-only** *percent-of-mbs*
  - **no high-prio-only**
  - **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
  - **no mbs**
  - **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}
  - **no packet-byte-offset**
  - **parent** {**root** | *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]
  - **no parent**
  - **percent-rate** *pir-percent* [**cir** *cir-percent*]
  - **no percent-rate**
  - **rate** {**max** | **kilobits-per-second**} [**cir** {**max** | **kilobits-per-second**}]
  - **no rate**
  - **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-priority-no-cir** | **offered-limited-profile-cir** | **offered-profile-cir** | **offered-priority-cir** | **offered-total-cir** | **offered-profile-capped-cir** | **offered-limited-capped-cir**}

— **no stat-mode**
— **queue** *queue-id* [**multipoint**] [*queue-type*] [*queue-mode*] **pool** *pool-name*
— **no queue** *queue-id*
— **adv-config-policy** ***policy-name***
— **no adv-config-policy**
— **burst-limit**
— **no burst-limit**
— **cbs** *size-in-kbytes*
— **no cbs**
— **high-prio-only** *percent*
— **no high-prio-only**
— **mbs** *size* [**bytes**|**kilobytes**]
— **no mbs**
— **parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
— **no parent**
— **percent-rate** *pir-percent* [**cir** *cir-percent*] [**port-limit**|**local-limit**]
— **percent-rate** *pir-percent* **police** [**port-limit**|**local-limit**]
— **no percent-rate**
— **rate** *pir-rate* [**cir** *cir-rate* | **police**]
— **no rate**
— **scope** {**exclusive** | **template**}
— **no scope**

Service Egress QoS Policy Commands

```
config
    — qos
        — [no] sap-egress policy-id | policy-name
            — description description-string
            — no description
            — dot1p dot1p-value [fc fc-name] [profile {in |out | use-de}]
            — no dot1p dot1p-value
            — dscp dscp-list [hsmda-counter-override counter-id] [fc fc-name] [profile {in | out}]
            — no dscp dscp-list
            — ethernet-ctag
            — no ethernet-ctag
            — fc fc-name
            — no fc fc-name
                — [no] de-mark [force de-value]
                — [no] de-mark-inner [force de-value]
                — [no] de-mark-outer [force de-value]
                — [no] dot1p {dot1p-value |in-profile dot1p-value out-profile dot1p-value}
                — [no] dot1p-inner {dot1p-value |in-profile dot1p-value out-profile dot1p-value}
                — [no] dot1p-outer {dot1p-value |in-profile dot1p-value out-profile dot1p-value}
                — dscp {dscp-name | in-profile dscp-name out-profile dscp-name}
                — no dscp
                — hsmda
                    — queue [1..8]
                    — no queue
                — policer policer-id [{[port-redirect-group-queue] [queue queue-id]} | {group queue-group-name [instance instance-id] [queue group-queue-id]}]
                — no policer
                — prec ip-prec-value [hsma-counter-override counter-id] [ fc fc-name] [profile {in | out}]
                — no prec ip-prec-value
                — queue queue-id [{group queue-group-name [instance instance-id]} | port-redirect-group-queue]
                — no queue
            — hsmda-queues
                — [no] low-burst-max-class class-id
                — packet-byte-offset {add add-bytes | subtract sub-bytes}
                — no packet-byte-offset
                — queue queue-id [port-redirect-group-queue]
                — no queue queue-id
                    — adaptation-rule [pir adaptation-rule]
                    — no adaptation-rule
                    — burst-limit size [bytes|kilobytes]
                    — no burst-limit
                    — mbs {[0..2625][kilobytes] | [0..2688000]bytes | default }
                    — no mbs
                    — rate pir-rate {max | kilobits-per-second}
                    — no rate
                    — slope-policy hsmda-slope-policy-name
```

```
                    — no slope-policy
                    — [no] wrr-weight weight
               — [no] wrr-policy wrr-policy-name
          — [no] ip-criteria
               — entry entry-id [create]
               — no entry entry-id
                    — action [hsmda-counter-override counter-id] [fc fc-name] [profile
                      {in | out}]
                    — no action
                    — description description string
                    — no description
                    — match [protocol protocol-id]
                    — no match
                         — dscp dscp-name
                         — no dscp
                         — dst-ip{ip-address/mask|ip-address ipv4-address-mask}
                         — no dst-ip
                         — dst-port {lt|gt|cq} dst-port-number
                         — dst-port range start end
                         — no dst-port
                         — fragment {true|false}
                         — no fragment
                         — src-ip {ip-address/mask|ip-address ipv4-address-mask}
                         — no src-ip
                         — src-port {lt|gt|eq} src-port-number
                         — src-port range start end
                         — no src-portport
               — renum old-entry-id new-entry-id
          — [no] ipv6-criteria
               — entry entry-id [create]
               — no entry entry-id
                    — action [hsmda-counter-override counter-id] [fc fc-name] [profile
                      {in | out}]
                    — no action
                    — description description string
                    — no description
                    — match [next-header next header]
                    — no match
                         — dscp dscp-name
                         — no dscp
                         — dst-ip {ipv6-address/prefix-length}
                         — no dst-ip
                         — dst-port {lt|gt|cq} dst-port-number
                         — dst-port range start end
                         — no dst-port
                         — src-ip {ipv6-address/prefix-length}
                         — no src-ip
                         — src-port {lt|gt|cq} dst-port-number
                         — src-port range start end
                         — no src-port
               — renum old-entry-number new-entry-number
          — parent-location {default|sla}
          — no parent-location
          — policer policer-id [group queue-group-name [queue group-queue-id]]
```

- **no policer**
  - **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
  - **no adaptation-rule**
  - **adv-config-policy** *adv-config-policy-name*
  - **no adv-config-policy**
  - **cbs** {*size* [**bytes** | **kilobytes**] | **default**}
  - **no cbs**
  - **description** *description string*
  - **no description**
  - **high-prio-only** *percent-of-mbs*
  - **no high-prio-only**
  - **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
  - **no mbs**
  - **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}
  - **no packet-byte-offset**
  - **parent** {**root** | *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]
  - **no parent**
  - **percent-rate** *percent-of-line-rate* [**cir** *percent-of-line-rate*]
  - **no percent-rate**
  - [**no**] **profile-capped**
  - [**no**] **profile-out-preserve**
  - **rate** {**max** | **kilobits-per-second**} [**cir** {**max** | **kilobits-per-second**}]
  - **no rate**
  - **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-profile-cir** | **offered-total-cir**}
  - **no stat-mode**
- **policy-name** *policy-name*
- **no policy-name**
- **prec** *ip-prec-value* [**hsmda-counter-override** *counter-id*] [**fc** *fc-name*] [**profile** {**in**|**out**}]
- **no prec** *ip-prec-value*
- **queue** *queue-id* [**multipoint**] [*queue-type*] [*queue-mode*] **pool** *pool-name*
- **no queue** *queue-id*
  - **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
  - **no adaptation-rule**
  - **adv-config-policy** *policy-name*
  - **no adv-config-policy**
  - **avg-frame-overhead** *percent*
  - **no avg-frame-overhead**
  - **burst-limit**
  - **no burst-limit**
  - **cbs** *size-in-kbytes*
  - **no cbs**
  - **high-prio-only** *percent*
  - **no high-prio-only**
  - **mbs** *size* [**bytes** | **kilobytes**]
  - **no mbs**
  - **parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
  - **no parent**
  - **percent-rate** *pir-percent* [**cir** *cir-percent*] [**port-limit** | **local-limit**]
  - **no percent-rate**
  - [**no**] **pool** *pool-name*
  - **port-parent** [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]

```
                — no port-parent
                — rate pir-rate [cir cir-rate]
                — no rate
                — xp-specific
                        — packet-byte-offset add add-bytes
                        — packet-byte-offset subtract sub-bytes
                        — no packet-byte-offset
                        — wred-queue [policy slope-policy-name]
                        — no wred-queue
— scope {exclusive | template}
— no scope
                —
```

Operational Commands

config
- **qos**
  - **copy sap-egress** *src-pol dst-pol* [**overwrite**]
  - **copy sap-ingress** *src-pol dst-pol* [**overwrite**]

Show Commands

show
- **qos**
  - **sap-ingress** *policy-id* [**detail**]
  - **sap-egress**[*policy-id*] [**association** | **match-criteria** | **hsmda** | **detail**]

Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
no description

Context config>qos>sap-egress
config>qos>sap-egress>ip-criteria>entry
config>qos>sap-ingress
config>qos>sap-ingress>ip-criteria>entry
config>qos>sap-ingress>ipv6-criteria>entry
config>qos>sap-ingress>mac-criteria>entry

Description This command creates a text description stored in the configuration file for a configuration context.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax **copy sap-egress** *src-pol dst-pol* [**overwrite**]
copy sap-ingress *src-pol dst-pol* [**overwrite**]

Context config>qos

Description This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id.

The **copy** command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters **sap-egress** *src-pol dst-pol* **—** Indicates that the source policy ID and the destination policy ID are sap-egress policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

Values 1 — 65535

sap-ingress *src-pol dst-pol* **—** Indicates that the source policy ID and the destination policy ID are SAP ingress policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

Values 1 — 65535

overwrite — Specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, an error will occur if the destination policy ID exists.

```
SR>config>qos# copy sap-egress 1 1010
MINOR: CLI Destination "1010" exists use {overwrite}.
SR>config>qos# copy sap-egress 1 1010 overwrite
```

renum

Syntax **renum** *old-entry-number new-entry-number*

Context config>qos>sap-ingress>ip-criteria
config>qos>sap-ingress>ipv6-criteria
config>qos>sap-egress>ipv6-criteria
config>qos>sap-ingress>mac-criteria

Description This command renumbers existing QoS policy criteria entries to properly sequence policy entries.

This can be required in some cases since the router exits when the first match is found and executes the actions in accordance with the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.

Parameters *old-entry-number* — Enter the entry number of an existing entry.

Default none

Values 1 — 65535

new-entry-number — Enter the new entry-number to be assigned to the old entry.

Default none

Values 1 — 65535

Service Ingress QoS Policy Commands

sap-ingress

Syntax [**no**] **sap-ingress** *policy-id* | *policy-name*

Context config>qos

Description This command is used to create or edit the ingress policy. The ingress policy defines the Service Level Agreement (SLA) enforcement service packets receive as they ingress a SAP. SLA enforcement is accomplished through the definition of queues that have Forwarding Class (FC), Committed Information Rate (CIR), Peak Information Rate (PIR), Committed Burst Size (CBS), and Maximum Burst Size (MBS) characteristics. The simplest policy defines a single queue that all ingress traffic flows through. Complex policies have multiple queues combined with specific IP or MAC match criteria that indicate which queue a packet will flow though.

Policies in effect are templates that can be applied to multiple services as long as the **scope** of the policy is template. Queues defined in the policy are not instantiated until a policy is applied to a service SAP.

A SAP ingress policy is considered incomplete if it does not include definition of at least one queue and does not specify the default action. The OS does not allow incomplete SAP ingress policies to be applied to services.

SAP ingress policies can be defined with either IP headers as the match criteria or MAC headers as the match criteria. The IP and MAC criteria are mutually exclusive and cannot be part of the same SAP ingress policy.

It is possible that a SAP ingress policy will include the **dscp** map command, the **dot1p** map command and an IP or MAC match criteria. When multiple matches occur for the traffic, the order of precedence will be used to arrive at the final action. The order of precedence is as follows:

1. 802.1p bits
2. DSCP
3. IP Quintuple or MAC headers

The SAP ingress policy with *policy-id* 1 is a system-defined policy applied to services when no other policy is explicitly specified. The system SAP ingress policy can be modified but not deleted. The **no sap-ingress** command restores the factory default settings when used on *policy-id* 1. The default SAP ingress policy defines one queue associated with the best effort (`be`) forwarding class, with CIR of zero and PIR of line rate.

Any changes made to the existing policy, using any of the sub-commands are applied immediately to all services where this policy is applied. For this reason, when many changes are required on a policy, it is recommended that the policy be copied to a work area policy ID. That work-in-progress policy can be modified until complete and then written over the original policy-id. Use the **config qos copy** command to maintain policies in this manner.

The **no sap-ingress** *policy-id* command deletes the SAP ingress policy. A policy cannot be deleted until it is removed from all services where it is applied. The system default SAP ingress policy is a special case; the **no** command restores the factory defaults to policy-id 1.

Parameters *policy-id* — The *policy-id* uniquely identifies the policy.

Values 1 — 65535

policy-name — The *policy-name* uniquely identifies the policy.

Values Valid names consist of any string up to 64 characters long. Policies must first be created with a policy-id, after which a policy-name can be assigned and used as an alias to reference the policy during configuration changes. Policy names may not begin with a number (0-9) or the underscore "_" character (e.g. _myPolicy). "Default" or "default" can not be used as policy names. Saved configurations and display output from the "info" and most "show" commands will show the policy-id (not the policy-name) where the policies are referenced.

policy-name

Syntax **policy-name** *policy-name*
no policy-name

Context cconfig>qos>sap-ingress
config>qos>sap-egress

Description Policies must first be created with a policy-id, after which a policy-name can be assigned and used as an alias to reference the policy during configuration changes. Saved configurations and display output from the **info** and most **show** commands will show the policy-id (not the policy-name) where the policies are referenced.

Default no policy-name

Parameters *policy-name* — Policy names may not begin with a number (0-9) or the underscore "_" character (e.g. _myPolicy). "Default" or "default" cannot be used as policy names. Specify a character string 64 characters or less.

scope

Syntax **scope** {**exclusive** | **template**}
no scope

Context config>qos>sap-ingress *policy-id*

Description This command configures the Service Ingress QoS policy scope as exclusive or template.

The policy's scope cannot be changed if the policy is applied to a service.

The **no** form of this command sets the scope of the policy to the default of **template**.

Default template

Parameters **exclusive —** When the scope of a policy is defined as exclusive, the policy can only be applied to one SAP. If a policy with an exclusive scope is assigned to a second SAP an error message is generated. If the

policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.

The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1.

template — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs on the router.

Default QoS policies are configured with template scopes. An error is generated when the template scope parameter to exclusive scope on default policies is modified.

default-fc

Syntax **default-fc** *fc-name*

Context config>qos>sap-ingress

Description This command configures the default forwarding class for the policy. In the event that an ingress packet does not match a higher priority (more explicit) classification command, the default forwarding class or sub-class will be associated with the packet. Unless overridden by an explicit forwarding class classification rule, all packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class. Optionally, the default ingress enqueuing priority for the traffic can be overridden as well.

The default forwarding class is best effort (`be`). The **default-fc** settings are displayed in the **show configuration** and **save** output regardless of inclusion of the **detail** keyword.

Context be

Parameters *fc-name* **—** Specify the forwarding class name for the queue. The value given for *fc-name* must be one of the predefined forwarding classes in the system.

The sub-class-name parameter is optional and used with the fc-name parameter to define a preexisting sub-class. The fc-name and sub-class-name parameters must be separated by a period (dot). If sub-class-name does not exist in the context of fc -name, an error will occur. If sub-class-name is removed using the no fc fc-name.sub-class-name force command, the default-fc command will automatically drop the sub-class-name and only use fc-name (the parent forwarding class for the sub-class) as the forwarding class.

Values

fc:	class[.sub-class]	
	class:	be, l2, af, l1, h2, ef, h1, nc
	sub-class:	29 characters max

Default None (Each sub-class-name must be explicitly defined)

default-priority

Syntax **default-priority** {**high** | **low**}

Context config>qos>sap-ingress

Description This command configures the default enqueuing priority for all packets received on an ingress SAP using this policy. To change the default priority for the policy, the **fc-name** must be defined whether it is being changed or not.

Default low

Parameters **high —** Setting the enqueuing parameter to **high** for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

low — Setting the enqueuing parameter to **low** for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

fc

Syntax [**no**] **fc** *fc-name*

Context config>qos>sap-ingress

Description The **fc** command creates a class or sub-class instance of the forwarding class fc-name. Once the *fc-name* is created, classification actions can be applied and the sub-class can be used in match classification criteria. Attempting to use an undefined sub-class in a classification command will result in an execution error and the command will fail.

The **no** form of the command removes all the explicit queue mappings for *fc-name* forwarding types. The queue mappings revert to the default queues for *fc-name*. To successfully remove a sub-class, all associations with the sub-class in the classification commands within the policy must first be removed or diverted to another forwarding class or sub-class.

Parameters *fc-name* — The parameter sub-class-name is optional and must be defined using a dot separated notation with a preceding valid system-wide forwarding class name. Creating a sub-class follows normal naming conventions. Up to sixteen ASCII characters may be used. If the same sub-name is used with two or more forwarding class names, each is considered a different instance of sub-class. A sub-class must always be specified with its preceding forwarding class name. When a forwarding class is created or specified without the optional sub-class, the parent forwarding class is assumed.

Within the SAP ingress QoS policy, up to 56 sub classes may be created. Each of the 56 sub-classes may be created within any of the eight parental forwarding classes. Once the limit of 56 is reached, any further sub-class creations will fail and the sub-class will not exist.

Successfully creating a sub-class places the CLI within the context of the sub-class for further sub-class parameter definitions. Within the sub-class context, commands may be executed that define sub-class

priority (within the parent forwarding class queue mapping), sub-class color aware profile settings, sub-class in-profile and out-of-profile precedence or DSCP markings.

The sub-class-name parameter is optional and used with the fc-name parameter to define a preexisting sub-class. The fc-name and sub-class-name parameters must be separated by a period (dot). If sub-class-name does not exist in the context of fc -name, an error will occur. If sub-class-name is removed using the no fc fc-name.sub-class-name force command, the default-fc command will automatically drop the sub-class-name and only use fc-name (the parent forwarding class for the sub-class) as the forwarding class.

Values fc: class[.sub-class]
class: be, l2, af, l1, h2, ef, h1, nc
sub-class: 29 characters max

Default None (Each sub-class-name must be explicitly defined)

policer

Syntax **policer** *policer-id* [**fp-redirect-group**]
no policer

Context config>qos>sap-ingress>fc

Description Within a sap-ingress QoS policy forwarding class context, the **policer** command is used to map packets that match the forwarding class and are considered unicast in nature to the specified policer-id. The specified policer-id must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination. If ingress forwarding logic has resolved a unicast destination (the packet does not need to be sent to multiple destinations), it is considered to be a unicast packet and will be mapped to either an ingress queue (using the **queue** *queue-id* or **queue** *queue-id* **group** *ingress-queue-group* commands) or an ingress policer (**policer** *policer-id*). The **queue** and **policer** commands within the forwarding class context are mutually exclusive. By default, the unicast forwarding type is mapped to the SAP ingress default queue (queue 1). If the **policer** *policer-id* command is executed, any previous policer mapping or queue mapping for the unicast forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast, unknown or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or ingress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class forwarding type mapping will fail.

When the unicast forwarding type within a forwarding class is mapped to a policer, the unicast packets classified to the sub-classes within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the unicast forwarding type within the forwarding class to the default queue. If all forwarding class forwarding types had been removed from the default queue, the queue will not exist on the SAPs or subscribers associated with the QoS policy and the **no policer** command will cause the system to attempt to create the default queue on each object. If the system cannot create the default queue in each instance, the **no policer** command will fail and the unicast forwarding type within the forwarding class will continue its mapping to the existing policer-id. If the **no policer** command results in a policer without any current mappings, the policer will be removed from the

SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

Parameters *policer-id* — When the forwarding class **policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.

Values 1—32

Default None

fp-redirect-group — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

broadcast-policer

Syntax **broadcast-policer** *policer-id* [**fp-redirect-group**]
no broadcast-policer

Context config>qos>sap-ingress>fc

Description Within a **sap-ingress** QoS policy forwarding class context, the **broadcast-policer** command is used to map packets that match the forwarding class and are considered broadcast in nature to the specified policer-id. The specified policer-id must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination based on the ingress service type and the service instance forwarding records. If the service type is VPLS and the destination MAC address is the broadcast address (ff:ff:ff:ff:ff:ff), the packet is classified into the broadcast forwarding type.

Broadcast forwarding type packets are mapped to either an ingress multipoint queue (using the **broadcast** *queue-id* or **broadcast** *queue-id* **group** *ingress-queue-group* commands) or an ingress policer (**broadcast-policer** *policer-id*). The **broadcast** and **broadcast-policer** commands within the forwarding class context are mutually exclusive. By default, the broadcast forwarding type is mapped to the SAP ingress default multipoint queue. If the **broadcast-policer** *policer-id* command is executed, any previous policer mapping or queue mapping for the broadcast forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast, unknown or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or ingress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class forwarding type mapping will fail.

The **broadcast-policer** command is ignored for instances of the policer applied to SAPs or subscribers where broadcast packets are not supported.

When the broadcast forwarding type within a forwarding class is mapped to a policer, the broadcast packets classified to the sub-classes within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the broadcast forwarding type within the forwarding class to the default multipoint queue. If all forwarding class forwarding types had been removed from the default multipoint queue, the queue will not exist on the SAPs or subscribers associated with the

QoS policy and the **no broadcast-policer** command will cause the system to attempt to create the default multipoint queue on each object. If the system cannot create the queue on each instance, the **no broadcast-policer** command will fail and the broadcast forwarding type within the forwarding class will continue its mapping to the existing policer-id. If the **no broadcast-policer** command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

Parameters *policer-id* — When the forwarding class **broadcast-policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the sap-ingress QoS policy.

Values 1—32

Default None

fp-redirect-group — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

multicast-policer

Syntax **multicast-policer** *policer-id* [**fp-redirect-group**]
no multicast-policer

Context config>qos>sap-ingress>fc

Description Within a **sap-ingress** QoS policy forwarding class context, the **multicast-policer** command is used to map packets that match the forwarding class and are considered multicast in nature to the specified policer-id. The specified policer-id must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination based on the ingress service type and the service instance forwarding records. Two basic types of services support multicast packets; routed services (IES and VPRN) and L2 multipoint services (VPLS, I-VPLS and B-VPLS). For the routed service types, a multicast packet is destined to an IPv4 or IPv6 multicast address. For the L2 multipoint services, a multicast packet is a packet destined to a multicast MAC address (multicast bit set in the destination MAC address but not the ff:ff:ff:ff:ff:ff broadcast address). The VPLS services also support two other multipoint forwarding types (broadcast and unknown) which are considered separate from the multicast forwarding type.

If ingress forwarding logic has resolved a packet to the multicast forwarding type within the forwarding class, it will be mapped to either an ingress multipoint queue (using the **multicast** *queue-id* or **multicast** *queue-id* **group** *ingress-queue-group* commands) or an ingress policer (**multicast-policer** *policer-id*). The **multicast** and **multicast-policer** commands within the forwarding class context are mutually exclusive. By default, the multicast forwarding type is mapped to the SAP ingress default multipoint queue. If the **multicast-policer** *policer-id* command is executed, any previous policer mapping or queue mapping for the multicast forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast, unknown or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or ingress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class forwarding type mapping will fail.

The multicast-policer command is ignored for instances of the policer applied to SAPs subscribers where broadcast packets are not supported.

When the multicast forwarding type within a forwarding class is mapped to a policer, the multicast packets classified to the sub-classes within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the multicast forwarding type within the forwarding class to the default multipoint queue. If all forwarding class forwarding types had been removed from the default multipoint queue, the queue will not exist on the SAPs subscribers associated with the QoS policy and the no multicast-policer command will cause the system to attempt to create the default multipoint queue on each object. If the system cannot create the queue on each instance, the no multicast-policer command will fail and the multicast forwarding type within the forwarding class will continue its mapping to the existing policer-id. If the no multicast-policer command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

Parameters *policer-id* — When the forwarding class **multicast-policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.

Values 1—32

Default None

fp-redirect-group — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

unknown-policer

Syntax **unknown-policer** *policer-id* **[fp-redirect-group]**
no unknown-policer

Context config>qos>sap-ingress>fc

Description Within a **sap-ingress** QoS policy forwarding class context, the **unknown-policer** command is used to map packets that match the forwarding class and are considered unknown in nature to the specified policer-id. The specified policer-id must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination based on the ingress service type and the service instance forwarding records. If the service type is VPLS and the destination MAC address is unicast but the MAC has not been learned and populated within the VPLS services FDB, the packet is classified into the unknown forwarding type.

Unknown forwarding type packets are mapped to either an ingress multipoint queue (using the **unknown** *queue-id* or **unknown** *queue-id* **group** *ingress-queue-group* commands) or an ingress policer (**unknown-policer** *policer-id*). The **unknown** and **unknown-policer** commands within the forwarding class context are mutually exclusive. By default, the unknown forwarding type is mapped to the SAP ingress default multipoint queue. If the **unknown-policer** *policer-id* command is executed, any previous policer mapping or queue mapping for the unknown forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast,

unknown or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or ingress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class forwarding type mapping will fail.

The **unknown-policer** command is ignored for instances of the policer applied to SAPs or subscribers where unknown packets are not supported.

When the unknown forwarding type within a forwarding class is mapped to a policer, the unknown packets classified to the sub-classes within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the unknown forwarding type within the forwarding class to the default multipoint queue. If all forwarding class forwarding types had been removed from the default multipoint queue, the queue will not exist on the SAPs or subscriber associated with the QoS policy and the no broadcast-policer command will cause the system to attempt to create the default multipoint queue on each object. If the system cannot create the queue on each instance, the no unknown-policer command will fail and the unknown forwarding type within the forwarding class will continue its mapping to the existing policer-id. If the no unknown-policer command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

Parameters *policer-id* — When the forwarding class **unknown-policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.

Values 1—32

Default None

fp-redirect-group — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

dot1p

Syntax **dot1p** *dot1p-value* [**fc** *fc-name*] [**profile** {**in |out** | **use-de**}]
no dot1p *dot1p-value*

Context config>qos>sap-ingress

Description This command explicitly sets the forwarding class or sub-class or enqueuing priority when a packet is marked with a *dot1p-priority* specified. Adding a dot1p rule on the policy forces packets that match the *dot1p-priority* specified to override the forwarding class and enqueuing priority based on the parameters included in the dot1p rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueuing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

The *dot1p-priority* is derived from the most significant three bits in the IEEE 802.1Q or IEEE 802.1P header. The three dot1p bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop Quality-of-Service (QoS) behavior.

The **no** form of this command removes the explicit dot1p classification rule from the SAP ingress policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

Parameters *dot1p-value* — This value is a required parameter that specifies the unique IEEE 802.1P value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

Values 0 — 7

fc *fc-name* **—** The value given for the fc-name parameter must be one of the predefined forwarding classes in the system. Specifying the fc-name is optional. When a packet matches the rule, the forwarding class is only overridden when the fc fc-name parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

Default None

priority {in|out|use-de} — All frames that are assigned to this forwarding class will be considered in or out of profile based on this command or to use the default. In case of congestion, the in- profile frames are preferentially queued over the out-of-profile frames

Values **in** — All frames are treated as in-profile.

out — All frames are treated as out of profile.

use-de — The profile of all frames is set according to the DEI bit.

dscp

Syntax **dscp** *dscp-name* [*dscp-name...*(*upto 8 max*)] **fc** *fc-name* [**priority** {**low** | **high**}]
no dscp *dscp-name* [*dscp-name...*(*upto 8 max*)]

Context config>qos>sap-ingress

Description This command explicitly sets the forwarding class or subclass or enqueuing priority when a packet is marked with the DiffServ Code Point (DSCP) value contained in the *dscp-name*. A list of up to 8 dscp-names can be entered on a single command. The lists of dscp-names within the configuration are managed by the system to ensure that each list does not exceed 8 names. Entering more than 8 dscp-names with the same parameters (**fc**, **priority**) will result in multiple lists being created. Conversely, multiple lists with the same parameters (fc, priority) are merged and the lists repacked to a maximum of 8 per list if dscp-names are removed or the parameters changed so the multiple lists use the same parameters. Also, if a subset of a list is entered with different parameters then a new list will be created for the subset. Note that when the list is stored in the configuration, the dscp-names are sorted by their DSCP value in ascending numerical order, consequently the order in the configuration may not be exactly what the user entered.

Adding a DSCP rule on the policy forces packets that match the DSCP value specified to override the forwarding class and enqueuing priority based on the parameters included in the DSCP rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueuing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

The DSCP value (referred to here by *dscp-name*) is derived from the most significant six bits in the IPv4

header ToS byte field (DSCP bits) or the Traffic Class field from the IPv6 header. If the packet does not have an IP header, dscp based matching is not performed. The six DSCP bits define 64 DSCP values used to map packets to per-hop Quality-of-Service (QoS) behavior. The most significant three bits in the IP header ToS byte field are also commonly used in a more traditional manner to specify an IP precedence value, causing an overlap between the precedence space and the DSCP space. Both IP precedence and DSCP classification rules are supported.

DSCP rules have a higher match priority than IP precedence rules and where a dscp-name DSCP value overlaps an ip-prec-value, the DSCP rule takes precedence.

The **no** form of the command removes the specified the *dscp-names* from the explicit DSCP classification rule in the SAP ingress policy. As *dscp-names* are removed, the system repacks the lists of dscp-names with the same parameters (up to 8 per list). As the **no** command does not have any additional parameters, it is possible to remove multiple *dscp-names* from multiple DSCP statements having different parameters with one command. If a *dscp-name* specified in a **no** command does not exist in any DSCP statement, then the command is aborted at that point with an error message displayed; any *dscp-names* in the list before the failed entry will be processed as normal but the processing will stop at the failed entry so that the remainder of the list is not processed.

Removing the *dscp-name* from the policy immediately removes the *dscp-name* on all ingress SAPs using the policy.

Parameters *dscp-name* — The DSCP name is a required parameter that specifies the unique IP header ToS byte DSCP bits value that will match the DSCP rule. If the command is executed multiple times with the same *dscp-name*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of 64 DSCP rules are allowed on a single policy and a maximum of 8 *dscp-names* can be specified in a single statement.

The specified name must exist as a *dscp-name*. SR OS software provides names for the well-known code points these can be shown using the command below:

```
A:PE# show qos dscp-table
============================================================
DSCP Mapping
============================================================
DSCP Name      DSCP Value     TOS (bin)      TOS (hex)
------------------------------------------------------------
be             0              0000 0000      00
cp1            1              0000 0100      04
cp2            2              0000 1000      08
cp3            3              0000 1100      0C
cp4            4              0001 0000      10
cp5            5              0001 0100      14
cp6            6              0001 1000      18
cp7            7              0001 1100      1C
cs1            8              0010 0000      20
cp9            9              0010 0100      24
af11           10             0010 1000      28
cp11           11             0010 1100      2C
af12           12             0011 0000      30
cp13           13             0011 0100      34
af13           14             0011 1000      38
cp15           15             0011 1100      3C
cs2            16             0100 0000      40
```

```
cp17            17              0100 0100       44
af21            18              0100 1000       48
cp19            19              0100 1100       4C
af22            20              0101 0000       50
cp21            21              0101 0100       54
af23            22              0101 1000       58
cp23            23              0101 1100       5C
cs3             24              0110 0000       60
cp25            25              0110 0100       64
af31            26              0110 1000       68
cp27            27              0110 1100       6C
af32            28              0111 0000       70
cp29            29              0111 0100       74
af33            30              0111 1000       78
cp31            31              0111 1100       7C
cs4             32              1000 0000       80
cp33            33              1000 0100       84
af41            34              1000 1000       88
cp35            35              1000 1100       8C
af42            36              1001 0000       90
cp37            37              1001 0100       94
af43            38              1001 1000       98
cp39            39              1001 1100       9C
cs5             40              1010 0000       A0
cp41            41              1010 0100       A4
cp42            42              1010 1000       A8
cp43            43              1010 1100       AC
cp44            44              1011 0000       B0
cp45            45              1011 0100       B4
ef              46              1011 1000       B8
cp47            47              1011 1100       BC
nc1             48              1100 0000       C0
cp49            49              1100 0100       C4
cp50            50              1100 1000       C8
cp51            51              1100 1100       CC
cp52            52              1101 0000       D0
cp53            53              1101 0100       D4
cp54            54              1101 1000       D8
cp55            55              1101 1100       DC
nc2             56              1110 0000       E0
cp57            57              1110 0100       E4
cp58            58              1110 1000       E8
cp59            59              1110 1100       EC
cp60            60              1111 0000       F0
cp61            61              1111 0100       F4
cp62            62              1111 1000       F8
cp63            63              1111 1100       FC
===============================================================
```

fc *fc-name* **—** The value given for *fc-name* must be one of the predefined forwarding classes in the system. Specifying the *fc-name* is optional. When a packet matches the rule the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The sub-class-name parameter is optional and used with the fc-name parameter to define a preexisting sub-class. The fc-name and sub-class-name parameters must be separated by a period (dot). If sub-

class-name does not exist in the context of fc -name, an error will occur. If sub-class-name is removed using the no fc fc-name.sub-class-name force command, the default-fc command will automatically drop the sub-class-name and only use fc-name (the parent forwarding class for the sub-class) as the forwarding class.

Values fc: class[.sub-class]
class: be, l2, af, l1, h2, ef, h1, nc
sub-class: 29 characters max

Default Inherit (When **fc** *fc-name* is not defined, the rule preserves the previous forwarding class of the packet.)

priority — This parameter overrides the default enqueuing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule the enqueuing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueuing priority is inherited based on previous rule matches.

Default low priority

high — This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **high** for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. Once the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Default low priority

low — This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **low** for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Default Inherit (When **priority** is not defined, the rule preserves the previous enqueuing priority of the packet.)

dscp

Syntax **dscp** *dscp-name* [*dscp-name...*(*upto 8 max*)] [**fc** *fc-name*] [**hsmda-counter-override** *counter-id*] [**profile** {**in** | **out**}]
no dscp *dscp-name* [*dscp-name...*(*upto 8 max*)]

Context config>qos>sap-egress

Description This command defines IP Differentiated Services Code Point (DSCP) names that must be matched to perform the associated reclassification actions. The specified name must exist as a *dscp-name*. SR OS software provides names for the well-known code points. A list of up to 8 *dscp-names* can be entered on a single command. The lists of *dscp-names* within the configuration are managed by the system to ensure that each list does not exceed 8 names. Entering more than 8 *dscp-names* with the same parameters (fc, hsmda-counter-override, priority) will result in multiple lists being created. Conversely, multiple lists with the same parameters (fc, hsmda-counter-override, priority) are merged and the lists repacked to a maximum of 8 per list if *dscp-names* are removed or the parameters changed so the multiple lists use the same parameters. Also, if a subset of a list is entered with different parameters then a new list will be created for the subset.

Note that when the list is stored in the configuration, the *dscp-names* are sorted by their DSCP value in ascending numerical order, consequently the order in the configuration may not be exactly what the user entered.

If an egress packet on the SAP matches an IP DSCP value corresponding to a specified *dscp-name*, the forwarding class, profile or HSMDA egress queue accounting behavior may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions. The default behavior for HSMDA queue accounting is to use the counters associated with the queue to which the packet is mapped. Matching a DSCP based reclassification rule will override all IP precedence based reclassification rule actions.

The IP DSCP bits used to match against dscp reclassification rules come from the Type of Service (ToS) field within the IPv4 header or the Traffic Class field from the IPv6 header. If the packet does not have an IP header, dscp based matching is not performed.

The reclassification actions from a dscp reclassification rule may be overridden by an IP flow match event.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions. If an ip-criteria match occurs after the DSCP match, the new forwarding class may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new fc, the fc from the dscp match will be used.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior. If an ip-criteria match occurs after the DSCP match, the new profile may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new profile, the profile from the DSCP match will be used.

The **hsmda-counter-override** keyword is optional. When specified, and the egress SAP is created on an HSMDA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMDA uses each queues default queue counters for packets mapped to the queue. The hsmda-counter-override keyword is used to map the packet to an explicit exception counter. One of eight counters may be used. When the packet is mapped to an exception counter, the packet will not increment the queues discard or forwarding counters, instead the exception discard and forwarding counters will be used.

The DSCP-based counter override decision may be overwritten by an ip-criteria reclassification rule match if the higher priority classification rule has an hsmda-counter-override action defined.

The **no** form of the command removes the specified the *dscp-names* from the reclassification rule in the SAP egress QoS policy. As *dscp-names* are removed, the system repacks the lists of *dscp-names* with the same parameters (up to 8 per list). As the **no** command does not have any additional parameters, it is possible to remove multiple *dscp-names* from multiple DSCP statements having different parameters with one command. If a *dscp-name* specified in a **no** command does not exist in any DSCP statement then the command is aborted at that point with an error message displayed. Any *dscp-names* in the list before the failed entry will be processed as normal but the processing will stop at the failed entry so that the remainder of the list is not processed.

Parameters *dscp-name:* — The *dscp-name* parameter is required when defining a DSCP reclassification rule. The specified name must exist as a dscp-name. A maximum of 8 dscp-names can be specified in a single statement. SR OS software provides names for the well-known code points, these can be shown using the command below:

```
A:PE# show qos dscp-table
============================================================
DSCP Mapping
============================================================
DSCP Name      DSCP Value     TOS (bin)      TOS (hex)
------------------------------------------------------------
be             0              0000 0000      00
cp1            1              0000 0100      04
cp2            2              0000 1000      08
cp3            3              0000 1100      0C
cp4            4              0001 0000      10
cp5            5              0001 0100      14
cp6            6              0001 1000      18
cp7            7              0001 1100      1C
cs1            8              0010 0000      20
cp9            9              0010 0100      24
af11           10             0010 1000      28
cp11           11             0010 1100      2C
af12           12             0011 0000      30
cp13           13             0011 0100      34
af13           14             0011 1000      38
cp15           15             0011 1100      3C
cs2            16             0100 0000      40
cp17           17             0100 0100      44
af21           18             0100 1000      48
cp19           19             0100 1100      4C
af22           20             0101 0000      50
cp21           21             0101 0100      54
af23           22             0101 1000      58
cp23           23             0101 1100      5C
cs3            24             0110 0000      60
cp25           25             0110 0100      64
af31           26             0110 1000      68
cp27           27             0110 1100      6C
af32           28             0111 0000      70
cp29           29             0111 0100      74
af33           30             0111 1000      78
cp31           31             0111 1100      7C
cs4            32             1000 0000      80
cp33           33             1000 0100      84
af41           34             1000 1000      88
cp35           35             1000 1100      8C
af42           36             1001 0000      90
cp37           37             1001 0100      94
af43           38             1001 1000      98
cp39           39             1001 1100      9C
cs5            40             1010 0000      A0
cp41           41             1010 0100      A4
cp42           42             1010 1000      A8
cp43           43             1010 1100      AC
cp44           44             1011 0000      B0
cp45           45             1011 0100      B4
ef             46             1011 1000      B8
```

```
cp47            47              1011 1100       BC
nc1             48              1100 0000       C0
cp49            49              1100 0100       C4
cp50            50              1100 1000       C8
cp51            51              1100 1100       CC
cp52            52              1101 0000       D0
cp53            53              1101 0100       D4
cp54            54              1101 1000       D8
cp55            55              1101 1100       DC
nc2             56              1110 0000       E0
cp57            57              1110 0100       E4
cp58            58              1110 1000       E8
cp59            59              1110 1100       EC
cp60            60              1111 0000       F0
cp61            61              1111 0100       F4
cp62            62              1111 1000       F8
cp63            63              1111 1100       FC
===============================================================
```

fc *fc-name:* — The **fc** reclassification action is optional. When specified, packets matching the IP DSCP value corresponding to a specified *dscp-name* will be explicitly reclassified to the forwarding class specified as fc-name regardless of the ingress classification decision. The explicit forwarding class reclassification may be overwritten by an ip-criteria reclassification match. The **fc** name defined must be one of the eight forwarding classes supported by the system. To remove the forwarding class reclassification action for the specified DSCP value, the **dscp** command must be re-executed without the **fc** reclassification action defined.

Values be, l1, af, l2, h1, ef, h2 or nc

profile {**in** | **out**} **—** The profile reclassification action is optional. When specified, packets matching the IP DSCP value corresponding to a specified *dscp-name* will be explicitly reclassified to either in-profile or out-of-profile regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by an ip-criteria reclassification match. To remove the profile reclassification action for the specified *dscp-name*, the **dscp** command must be re-executed without the profile reclassification action defined.

in — The **in** parameter is mutually excusive to the out parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When in is specified, any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane. This value may be overwritten by an explicit profile action in an ip-criteria reclassification match.

out: — The **out** parameter is mutually excusive to the in parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When out is specified, any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane. This value may be overwritten by an explicit profile action in an ip-criteria reclassification match.

hsmda-counter-override *counter-id* — The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the IP precedence value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The hsmda-counter-override action may be overwritten by an ip-criteria reclassification rule match. The specified counter-id must be specified as an integer between 1 and 8. To

remove the ESMDA exception counter reclassification action for the specified DSCP value, the dscp command must be re-executed without the hsmda-counter-override reclassification action defined.

Values 1 —8

ethernet-ctag

Syntax [**no**] **ethernet-ctag**

Context config>qos>sap-egress

Description This command specifies that the top customer tag should be used for egress reclassification based on dot1p criteria. This command applies to all dot1p criteria configured in a given SAP egress QoS policy.

The no form of this command means that a service delimiting tag will be used for egress reclassification based on dot1p criteria.

Default noethernet-ctag

ip-criteria

Syntax [**no**] **ip-criteria**

Context config>qos>sap-egress

Description IP criteria-based SAP ingress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.

This command is used to enter the context to create or edit policy entries that specify IP criteria such as IP quintuple lookup or DiffServ code point.

The software implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. Once IP criteria entries are removed from a SAP ingress policy, the IP criteria is removed from all services where that policy is applied.

ip-criteria

Syntax [**no**] **ip-criteria**

Context config>qos>sap-egress

Description IP criteria-based SAP egress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.

This command is used to enter the context to create or edit policy entries that specify IP criteria such as IP quintuple lookup or DiffServ code point.

The OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. Once IP criteria entries are removed from a SAP ingress policy, the IP criteria is removed from all services where that policy is applied.

ipv6-criteria

Syntax [**no**] **ipv6-criteria**

Context config>qos>sap-ingress

Description IPv6 criteria-based SAP ingress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.

This command is used to enter the node to create or edit policy entries that specify IPv6 criteria such as IP quintuple lookup or DiffServ code point.

The OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. Once ipv6-criteria entries are removed from a SAP ingress policy, the ipv6-criteria is removed from all services where that policy is applied.

lsp-exp

Syntax **lsp-exp** *lsp-exp-value* [**fc** *fc-name*] [**priority** {**low**|**high**}] [**hsmda-counter-override** *counter-id*]
no lsp-exp *lsp-exp-value*

Context config>qos>sap-ingress

Description This command explicitly sets the forwarding class or sub-class enqueuing priority when a packet is marked with a MPLS EXP bits specified. Adding a lsp-exp rule on the policy forces packets that match the MPLS LSP EXP specified to override the forwarding class and enqueuing priority based on the parameters included in the lsp-exp rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueuing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy..

The *lsp-exp-value* is derived from the MPLS LSP EXP bits of the top label.

Multiple commands can be entered to define the association of some or all eight LSP EX bit values to the forwarding class.

The **no** form of this command removes the explicit lsp-exp classification rule from the SAP ingress policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

This command applies to Ethernet Layer 2 SAPs only.

Default none

Parameters *lsp-exp-value* — This value is a required parameter that specifies the unique MPLS LSP EXP value that will match the lsp-exp rule. If the command is executed multiple times with the same lsp-exp-value, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight lsp-exp rules are allowed on a single policy.

Values 0 — 7

fc *fc-name* **—** The value given for the fc-name parameter must be one of the predefined forwarding classes in the system. Specifying the fc-name is optional. When a packet matches the rule the forwarding class is only overridden when the fc fc-name parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The sub-class-name parameter is optional and used with the fc-name parameter to define a preexisting sub-class. The fc-name and sub-class-name parameters must be separated by a period (dot). If sub-class-name does not exist in the context of fc -name, an error will occur.

Values fc: class[.sub-class]
class: be, l2, af, l1, h2, ef, h1, nc
sub-class: 29 characters max

Default None (Each sub-class-name must be explicitly defined)

priority — The priority parameter is used to override the default enqueuing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule, the enqueuing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueuing priority is inherited based on previous rule matches.

high — The high parameter is used in conjunction with the priority parameter. Setting the enqueuing parameter to high for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

low — The low parameter is used in conjunction with the priority parameter. Setting the enqueuing parameter to low for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Default No override.

hsmda-counter-override *counter-id* **—** The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the MPLS EXP value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The specified counter-id must be specified as an integer between 1 and 8. To remove the HSMDA exception counter reclassification action for the specified lsp-exp-value, the lsp-exp command must be re-executed without the hsmda-counter-override reclassification action defined.

Values 1 — 8

mac-criteria

Syntax [**no**] **mac-criteria**

Context config>qos>sap-ingress

Description The **mac-criteria** based SAP ingress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.

This command is used to enter the node to create or edit policy entries that specify MAC criteria.

Router implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. Once mac-criteria entries are removed from a SAP ingress policy, the mac-criteria is removed from all services where that policy is applied.

policer

Syntax **policer** *policer-id* [**create**]
no policer *policer-id*

Context config>qos>sap-ingress

Description This command is used in the sap-ingress and sap-egress QoS policies to create, modify or delete a policer. Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The sap-ingress policy may have up to 32 policers (numbered 1 through 32) may be defined while the sap-egress QoS policy supports a maximum of 8 (numbered 1 through 8). While a policer may be defined within a QoS policy, it is not actually created on SAPs or subscribers associated with the policy until a forwarding class is mapped to the policer's ID.

All policers must be created within the QoS policies. A default policer is not created when a sap-ingress or sap-egress QoS policy is created.

Once a policer is created, the policer's metering rate and profiling rates may be defined as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues which have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

Once a policer is created, it cannot be deleted from the QoS policy unless any forwarding classes that are mapped to the policer are first moved to other policers or queues.

The system will allow a policer to be created on a SAP QoS policy regardless of the ability to support policers on objects where the policy is currently applied. The system only scans the current objects for policer support and sufficient resources to create the policer when a forwarding class is first mapped to the policer ID. If the policer cannot be created due to one or more instances of the policy not supporting policing or having insufficient resources to create the policer, the forwarding class mapping will fail.

The **no** form of this command is used to delete a policer from a sap-ingress or sap-egress QoS policy. The specified policer cannot currently have any forwarding class mappings for the removal of the policer to

succeed. It is not necessary to actually delete the policer ID for the policer instances to be removed from SAPs or subscriber associated with the QoS policy once all forwarding classes have been moved away from the policer. It is automatically deleted from each policing instance although it still appears in the QoS policy.

Parameters *policer-id* — The *policer-id* must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword requirements which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer's context for possible parameter modification.

Values 1—32

description

Syntax **description** *description string*
no description

Context config>qos>sap-ingress>policer

Description The **description** command is used to define an informational ASCII string associated with the policer control policy. The string value can be defined or changed at any time once the policy exists.

The **no** form of this command is used to remove an explicit description string from the policer.

Default **no description**

Parameters *description string* — The *description-string* parameter defines the ASCII description string for the policer control policy. The description string can be up to 80 characters. If the string contains spaces, it must be placed within beginning and ending double quotation marks. Beginning and ending quotation marks are not considered part of the description string. Only printable ASCII characters are allowed in the string. The sting does not need to be unique and may be repeated in the descriptions for other policer control policies or other objects. If the command is executed without the description-sting present, any existing description string will be unaffected.

Default None

adv-config-policy

Syntax [**no**] **adv-config-policy** *policy-name*

Context config>qos>sap-ingress>policer
config>qos>sap-egress>policer

Description This command specifies the advanced QoS policy. The advanced QoS policy contains only queue and policer child control parameters within a child-control node.

Once a policy is created, it may be applied to a queue or policer defined within a **sap-egress** or **sap-ingress** QoS policy. It may also be applied to a queue or policer defined within an ingress or **egress queue-group**

template. When a policy is currently associated with a QoS policy or template, the policy may be modified but not deleted (even in the event that the QoS policy or template is not in use).

The **no** form of this command removes the specified advanced policy.

Default None

Parameters *policy-name* — The name of the advanced QoS policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

adaptation-rule

Syntax **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
no adaptation-rule

Context config>qos>sap-ingress>policer

Description This command is used to define how the policer's configuration parameters are translated into the underlying hardware capabilities used to implement each policer instance. For instance, the configured rates for the policer need to be mapped to the timers and decrement granularity used by the hardware's leaky bucket functions that actually perform the traffic metering. If a rate is defined that cannot be exactly matched by the hardware, the adaptation-rule setting provides guidance for which hardware rate should be used.

The **max** keyword tells the system that the defined rate is the maximum rate that should be given to the policer. If the hardware cannot exactly match the given rate, the next lowest hardware supported rate is used.

The **min** keyword tells the system that the defined rate is the minimum rate that should be given to the policer. If the hardware cannot exactly match the given rate, the next highest hardware supported rate is used.

The **closest** keyword tells the system that the defined rate is the target rate for the policer. If the hardware cannot exactly match the given rate, the system will use the closest hardware supported rate compared to the target rate.

The hardware also needs to adapt the given mbs and cbs values into the PIR bucket violate threshold (discard) and the CIR bucket exceed threshold (out-of-profile). The hardware may not have an exact threshold match which it can use. In R8.0, the system treats the mbs and cbs values as minimum threshold values.

The **no** form of this command is used to return the policer's metering and profiling hardware adaptation rules to closest.

Parameters **pir** {**max** | **min** | **closest**} — When the optional **pir** parameter is specified, the **max**, **min** or **closest** keyword qualifier must follow.

max — The **max** keyword is used to inform the system that the metering rate defined for the policer is the maximum allowed rate. The system will choose a hardware supported rate that is closest but not exceeding the specified rate.

min — The **min** keyword is used to inform the system that the metering rate defined for the policer is the minimum allowed rate. The system will choose a hardware supported rate that is closest but not lower than the specified rate.

closest — The **closest** keyword is used to inform the system that the metering rate defined for the policer is the target rate. The system will choose a hardware supported rate that is closest to the specified rate.

Default closest

cir {**max** | **min** | **closest**} — When the optional **cir** parameter is specified, the **max**, **min** or **closest** keyword qualifier must follow.

max — The **max** keyword is used to inform the system that the profiling rate defined for the policer is the maximum allowed rate. The system will choose a hardware supported rate that is closest but not exceeding the specified rate.

min — The min keyword is used to inform the system that the profiling rate defined for the policer is the minimum allowed rate. The system will choose a hardware supported rate that is closest but not lower than the specified rate.

closest — The closest keyword is used to inform the system that the profiling rate defined for the policer is the target rate. The system will choose a hardware supported rate that is closest to the specified rate.

Default closest

Values

cbs

Syntax **cbs** {*size* [**bytes** | **kilobytes**] | **default**}
no cbs

Context config>qos>sap-ingress>policer
config>qos>qgrps>egr>qgrp>policer

Description This command is used to configure the policer's CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the policer. If the policer's forwarding rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.

The policer's **cbs** size defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The no form of this command returns the policer to its default CBS size.

Default 64 kilobytes when CIR = **max**, otherwise 10ms volume of traffic for a configured non zero/non max CIR.

Parameters *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **cbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

byte — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobyte — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Values 0 – 16777216 or **default**

Default **kilobyte**

high-prio-only

Syntax **high-prio-only** *percent-of-mbs*
no high-prio-only

Context config>qos>sap-ingress>policer

Description This command is used to configure the percentage of the policer's PIR leaky bucket's MBS (maximum burst size) that is reserved for high priority traffic. While the **mbs** value defines the policer's high priority violate threshold, the percentage value defined is applied to the **mbs** value to derive the bucket's low priority violate threshold. See the **mbs** command details for information on which types of traffic is associated with each violate threshold.

Default **high-prio-only 10**

Parameters *percent-of-mbs* — The *percent-of-mbs* parameter is required when specifying **high-prio-only** and is expressed as a percentage with granularity of 1,000th of a percent.

Values 0—100

Default 10

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
no mbs

Context config>qos>sap-ingress>policer

Description This command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low priority violate threshold. For ingress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold.

The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst

tolerance defined by **high-prio-only** is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.

The no form of this command returns the policer to its default MBS size.

Default 64 kilobytes when PIR = **max**, otherwise, 10ms volume of traffic for a configured non zero/non max PIR.

Parameters *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

byte — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobyte — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Values 0 – 16777216 or **default**

Default **kilobyte**

packet-byte-offset

Syntax **packet-byte-offset add** *add-bytes*
packet-byte-offset subtract *sub-bytes*
no packet-byte-offset

Context config>qos>sap-egress>queue>xp-specific

Description This command is used to modify the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** version of this command is used to remove per packet size modifications from the policer.

Parameters **add** *bytes* — The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 0 — 31

Default None

subtract *bytes* — The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When **subtract** is defined, the corresponding *bytes* parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 1 — 32

Default None

packet-byte-offset

Syntax **packet-byte-offset add** *add-bytes*
packet-byte-offset subtract *sub-bytes*
no packet-byte-offset

Context config>qos>sap-egress>queue>xp-specific

Description This command is used to modify the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** version of this command is used to remove per packet size modifications from the policer.

Parameters **add** *bytes* — The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 0 — 32

Default None

subtract *sub-bytes* — The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When **subtract** is defined, the corresponding *bytes* parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 0—64

Default None

parent

Syntax **parent** {**root** | *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]
no parent

Context config>qos>sap-ingress>policer

Description This command is used to create a child to parent mapping between each instance of the policer and either the **root** arbiter or a specific tiered **arbiter** on the object where the policy is applied. Defining a parent association for the policer causes the policer to compete for the parent policer's available bandwidth with other child policers mapped to the policer control hierarchy.

Policer control hierarchies may be created on SAPs or on a subscriber context. To create a policer control hierarchy on an ingress or egress SAP context, a **policer-control-policy** must be applied to the SAP. Once applied, the system will create a parent policer that is bandwidth limited by the policy's **max-rate** value under the root arbiter. The root arbiter in the policy also provides the information used to determine the various priority level discard-unfair and discard-all thresholds. Besides the root arbiter, the policy may also contain user defined tiered arbiters that provide arbitrary bandwidth control for subsets of child policers that are either directly or indirectly parented by the arbiter.

When the QoS policy containing the policer with a **parent** mapping to an arbiter name exists on the SAP, the system will scan the available arbiters on the SAP. If an arbiter exists with the appropriate name, the policer to arbiter association is created. If the specified arbiter does not exist either because a **policer-control-policy** is not currently applied to the SAP or the arbiter name does not exist within the applied policy, the policer is placed in an orphan state. Orphan policers operate as if they are not parented and are not subject to any bandwidth constraints other than their own PIR. When a policer enters the orphan state, it is flagged as operationally degraded due to the fact that it is not operating as intended and a trap is generated. Whenever a **policer-control-policy** is added to the SAP or the existing policy is modified, the SAP's policer's parenting configurations must be reevaluated. If an orphan policer becomes parented, the degraded flag should be cleared and a resulting trap should be generated.

For subscribers, the policer control hierarchy is created through the **policer-control-policy** applied to the **sub-profile** used by the subscriber. A unique policer control hierarchy is created for each subscriber associated with the **sub-profile**. The QoS policy containing the policer with the parenting command comes into play through the subscriber **sla-profile** which references the QoS policy. The combining of the **sub-profile** and the **sla-profile** at the subscriber level provides the system with the proper information to create the policer control hierarchy instance for the subscriber.

Executing the **parent** command will fail if:

- The policer's stat-mode in the QoS policy is set to no-stats
- A stat-mode no-stats override exists on an instance of the policer on a SAP or subscriber context

A policer with a **parent** command applied cannot be configured with **stat-mode no-stats** in either the QoS policy or as an override on an instance of the policer

Once a policer is successfully parented to an arbiter, the **parent** commands **level** and **weight** parameters are used to determine at what priority level and at which weight in the priority level that the child policer competes with other children (policers or other arbiters) for bandwidth.

The **no** form of this command is used to remove the parent association from all instances of the policer.

Parameters {**root** | *arbiter-name*} — When the **parent** command is executed, either the keyword **root** or an *arbiter-name* must be specified.

root — The **root** keyword specifies that the policer is intended to become a child to the **root** arbiter where an instance of the policer is created. If the **root** arbiter does not exist, the policer will be placed in the orphan state.

Default **root**

arbiter-name — The *arbiter-name* parameter specifies that the policer is intended to become a child to one of the tiered arbiters with the given arbiter-name where an instance of the policer is created. If the specified arbiter-name does not exist, the policer will be placed in the orphan state.

Default None

weight *weight-within-level* — The **weight** *weight-within-level* keyword and parameter are optional when executing the **parent** command. When **weight** is not specified, a default level of 1 is used in the parent arbiters priority level. When **weight** is specified, the *weight-within-level* parameter must be specified as an integer value from 1 through 100.

Default 1

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*]
no percent-rate

Context config>qos>sap-egress>policer
config>qos>sap-ingress>policer

Description The percent-rate command within the SAP ingress and egress QOS policy enables supports for a policer's PIR and CIR rate to be configured as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

This enables the same QOS policy to be used on SAPs on different FPs without needing to use SAP based policer overrides to modify a policer's rate to get the same relative performance from the policer.

If the parent arbiter rate changes after the policer is created, the policer's PIR and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a policer is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A policer's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

The **no** form of this command returns the queue to its default shaping rate and cir rate.

Parameters *pir-percent* — Specifies the policer's PIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values Percentage ranging from 0.01 to 100.00.

Default 100.00

cir cir-percent — The **cir** keyword is optional and when defined the required cir-percent CIR parameter expresses the policer's CIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values Percentage ranging from 0.00 to 100.00.

Default 100.00

profile-capped

Syntax [**no**] **profile-capped**

Context config>qos>sap-ingress>policer
config>qos>sap-egress>policer
config>qos>queue-group-templates>ingress>queue-group
config>qos>queue-group-templates>ingress>queue-group

Description Profile capped mode enforces an overall in-profile burst limit to the CIR bucket for ingress undefined, ingress explicit in-profile, egress soft-in-profile, and egress explicit in-profile packets. The default behavior when profile-capped mode is not enabled is to ignore the CIR output state when an explicit in-profile packet is handled by an ingress or egress policer.

The profile-capped mode makes two changes:

- At egress, soft-in-profile packets (packets received from ingress as in-profile) are treated the same as explicit in-profile (unless explicitly reclassified as out-of-profile) and have an initial policer state of in-profile.
- At both ingress and egress, any packet output from the policer with a non-conforming CIR state are treated as out-of-profile (out-of-profile state is ignored for initial in-profile packets when profile capped mode is not enabled)

Default no profile-capped

profile-out-preserve

Syntax [**no**] **profile-out-preserve**

Context config>qos>sap-egress>policer

Description This command specifies whether to preserve the color of offered out-of-profile traffic at sap-egress policer (profility of the packet can change based on egress CIR state).

When enabled, traffic determined as out-of-profile at ingress policer will be treated as out-of-profile at sap-egress policer.

rate

Syntax **rate** {**max** | **kilobits-per-second**} [**cir** {**max** | **kilobits-per-second**}]
no rate

Context config>qos>sap-ingress>policer

Description This command is used to configure the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on the each packets size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches it's exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs**, and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow based on the conforming or exceeding state from the CIR bucket.

When a packet is red neither the PIR or CIR bucket depths are incremented by the packets size. When the packet is yellow the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 Kbps (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command is used to restore the default metering and profiling rate to a policer.

Parameters {**max** | *kilobits-per-second*} — Specifying the keyword **max** or an explicit *kilobits-per-second* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to max.

Values **max** or 1—2000000000

cir {**max** | *kilobits-per-second*} — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *kilobits-per-second* parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 Kbps. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When **max** is specified, the maximum policer rate

used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CPIR used is equivalent to max.

Values **max** or 0—2000000000

stat-mode

Syntax **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-priority-no-cir** | **offered-limited-profile-cir** | **offered-profile-cir** | **offered-priority-cir** | **offered-total-cir | offered-profile-capped-cir | offered-limited-capped-cir** }
no stat mode

Context config>qos>sap-ingress>policer
config>qos>queue-group-templates>ingress>queue-group

Description This command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, high priority or low priority) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the large number of policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported which prevents any packet accounting, the use of the policer's **parent** command requires at the policer's **stat-mode** to be set at least to the **minimal** setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. Once a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the the total/allocated/free stats by using the **tools dump system-resources** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The default **stat-mode** when a policer is created within the policy is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's stat-mode setting to minimal. The command will fail if insufficient policer counter resources exist to implement minimal where the QoS policer is currently applied and has a forwarding class mapping.

no-stats — Counter resource allocation:0

The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using no-stats cannot be a child to a parent policer and the policer's parent command will fail.

When **collect-stats** is enabled, the lack of counters causes the system to generate the following statistics:

a. offered-in = 0

b. offered-out = 0

c.'discard-in = 0

d. discard-out = 0

e. forward-in =0

f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

minimal — Counter resource allocation:1

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (profile or priority) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters are used in the following manner:

1. offered = profile in/out, priority high/low

2. discarded = Same as 1

3. forwarded = Derived from 1 - 2

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out = 0

c. discard-in = 2

d. discard-out = 0

e. forward-in = 3

f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

With **minimal** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-in = 1

ii. offered-out = 0

iii. offered-undefined = 0

iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-profile-no-cir — Counter resource allocation:2

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when the policer is receiving only in-profile and out-of-profile pre-marked (and trusted) packets. It is expected that in this instance a CIR rate will not be defined since all packet are already pre-marked. This mode does not prevent the policer from receiving un-trusted (color undefined) nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-in = profile in

2. offered-out = profile out, priority high/low

3. dropped-in = Same as 1

4. dropped-out = Same as 2

5. forwarded-in = Derived from 1 - 3

6. forwarded-out = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out = 2

c. discard-in = 3

d. discard-out = 4

e. forward-in = 5

f. forward-out = 6

With **offered-profile-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-in = 1

ii. offered-out = 2

iii. offered-undefined = 0

iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-priority-no-cir — Counter resource allocation:2

The **offered-priority-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-priority-no-cir** mode is most useful when the policer is receiving only un-trusted packets and the ingress priority high and priority low classification options are being used without a CIR profiling rate defined. This mode does not prevent the policer from receiving trusted packets that are pre-marked in-profile or out-of-profile nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-high = profile in, priority high
2. offered-low = profile out, priority low
3. dropped-high = Same as 1
4. dropped-low = Same as 2
5. forwarded-high = Derived from 1 - 3
6. forwarded-low = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-high = 1
b. offered-low = 2
c. discard-high = 3
d. discard-low = 4
e. forward-high = 5
f. forward-low = 6

With **offered-priority-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high = 1
ii. offered-low = 2
iii. offered-undefined = 0
iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-limited-profile-cir — Counter resource allocation:3

The **offered-limitied-profile-cir** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-limited-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile (profile out but no profile in) traffic and un-trusted packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile packets.

The counters are used in the following manner:

1. offered-undefined-that-turned-green = profile in, priority high/low
2. offered-undefined-that-turned-yellow-or-red = priority high/low
3. offered-out-that-stayed-yellow-or-turned-red = profile out

4. dropped-undefined-that-turned-green = Same as 1

5. dropped-undefined-that-turned-yellow-or-red = Same as 2

6. dropped-out-that-turned-yellow-or-red = Same as 3

7. forwarded-undefined-that-turned-green = Derived from 1 - 4

8. forwarded-undefined-that-turned-yellow = Derived from 2 - 5

9. forwarded-out-that-turned-yellow = Derived from 3 - 6

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 0

b. offered-out = 1 + 2 + 3

c. discard-in = 0

d. discard-out = 4 + 5 + 6

e. forward-in = 7

f. forward-out = 8 + 9

With **offered-limited-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-in = 0

ii. offered-out = 3

iii. offered-undefined = 1 + 2

iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-profile-cir — Counter resource allocation:4

The **offered-profile-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile and in-profile traffic and is also receiving un-trusted packets that are being applied to a defined CIR profiling rate. This mode differs from **offered-limited-profile-cir** mode in that it expects both trusted in-profile and out-of-profile packets while still performing CIR profiling on packets with un-trusted markings. It is expected that in most cases where both trusted and un-trusted packets are received, the predominate case will not include trusted in-profile packets making the offered-limited-profile-cir accounting mode acceptable.

The counters are used in the following manner:

1. offered-in-that-stayed-green-or-turned-red = profile in

2. offered-undefined-that-turned-green = priority high/low

3. offered-undefined-that-turned-yellow-or-red = priority high/low

4. offered-out-that-stayed-yellow-or-turned-red = profile out

5. dropped-in-that-stayed-green-or-turned-red = Same as 1

6. dropped-undefined-that-turned-green = Same as 2

7. dropped-undefined-that-turned-yellow-or-red = Same as 3

8. dropped-out-that-turned-yellow-or-red = Same as 4

9. forwarded-in-that-stayed-green = Derived from 1 - 5

10. forwarded-undefined-that-turned-green = Derived from 2 - 6

11. forwarded-undefined-that-turned-yellow = Derived from 3 - 7

12. forwarded-out-that-turned-yellow = Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out = 2 + 3 + 4

c. discard-in = 5 + 6

d. discard-out = 7 + 8

e. forward-in = 9 + 10

f. forward-out = 11 + 12

With **offered-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high = 1

ii. offered-low = 4

iii. offered-undefined = 2 + 3

iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-priority-cir — Counter resource allocation:4

The **offered-priority-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-priority-cir** mode is most useful when the policer is receiving only un-trusted packets that are being classified as high priority or low priority and are being applied to a defined CIR profiling rate. This mode differs from **offered-profile-cir** mode in that it does not expect trusted in-profile and out-of-profile packets but does not exclude the ability of the policer to receive them.

The counters are used in the following manner:

1. offered-high-that-turned-green = profile in, priority high

2. offered-high-that-turned-yellow-or-red = profile in, priority high

3. offered-low-that-turned-green = profile out, priority low

4. offered-low-that-turned-yellow-or-red = profile out, priority low

5. dropped-high-that-turned-green = Same as 1

6. dropped-high-that-turned-yellow-or-red = Same as 2

7. dropped-low-that-turned-green = Same as 3

8. dropped-low-that-turned-yellow-or-red = Same as 4

9. forwarded-high-that-turned-green = Derived from 1 - 5

10. forwarded-high-that-turned-yellow = Derived from 2 - 6

11. forwarded-low-that-turned-green = Derived from 3 - 7

12. forwarded-low-that-turned-yellow = Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-high = 1 + 2

b. offered-low = 3 + 4

c. discard-in = 5 + 7

d. discard-out = 6 + 8

e. forward-in = 9 + 11

f. forward-out = 10 + 12

With **offered-priority-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high = 1 + 2

ii. offered-low = 3 + 4

iii. offered-undefined = 0

iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-total-cir — Counter resource allocation:2

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high and low priority classifications are not being used on the un-trusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the un-trusted packets.

The counters are used in the following manner:

1. offered-that-turned-green = profile in/out, priority high/low

2. offered- that-turned-yellow-or-red = profile in/out, priority high/low

3. dropped-offered-that-turned-green = Same as 1

4. dropped-offered-that-turned-yellow-or-red = Same as 2

5. forwarded-offered-that-turned-green = Derived from 1 - 3

6. forwarded-offered-that-turned-yellow = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1 + 2

b. offered-out = 0

c. discard-in = 3

d. discard-out = 4

e. forward-in = 5

f. forward-out = 6

Counter 0 indicates that the accounting statistic returns a value of zero.

With **offered-total-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high = 1 + 2

ii. offered-low = 0

iii. offered-undefined = 0

iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-profile-capped-cir — Counter resource allocation:2

When **offered-profile-capped-cir** is defined, the system creates five offered-output counters in the forwarding plane and five discard counters in the traffic manager.

The **offered-profile-capped-cir** mode is similar to the offered-profile-cir mode except that it includes support for profile in and **soft-in-profile** that may be output as 'out-of-profile' due to enabling **profile-capped** mode on the ingress policer.

The impact of using **offered-profile-capped-cir** stat-mode while **profile-capped** mode is disabled are that one of the counting resources in the forwarding plane and traffic manager will not be used and soft-in-profile will be treated as 'offered-in' instead of 'offered-undefined'.

The counters are used in the following manner:

1. 'offered-in-that-stayed-green' = profile in, soft-in-profile
2. 'offered-in-that-turned-yellow-or-red' = profile in, soft-in-profile
3. 'offered-soft-out-that-turned-green = soft-out-of-profile
4. 'offered-soft-out- that-turned-yellow-or-red' = soft-out-of-profile
5. 'offered-out-that-turned-yellow-or-red' = profile out
6. 'dropped-in-that-stayed-green' = Same as 1
7. 'dropped-in-that-turned-yellow-or-red' = Same as 2
8. 'dropped-soft-out-that-turned-green' = Same as 3
9. 'dropped-soft-out-that-turned-yellow-or-red' = Same as 4
10. 'dropped-out-that-turned-yellow-or-red' = Same as 5

11. ‘forwarded-in-that-stayed-green’ = Derived from 1 - 6
12. ‘forwarded-in-that-turned-yellow’ = Derived from 2 - 7
13. ‘forwarded-soft-out-that-turned-green’ = Derived from 3 - 8
14. ‘forwarded-soft-out-that-turned-yellow’ = Derived from 4 - 9
15. ‘forwarded-out-that-turned-yellow’ = Derived from 5 - 10

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

a. ‘offered-undefined’ = 3 + 4
b. ‘offered-in’ = 1 + 2
c. ‘offered-out’ = 5
d. ‘discard-in’ = 6 + 8
e. ‘discard-out’ = 7 + 9 + 10
f. ‘forward-in’ = 11 + 13
g. ‘forward-out’ = 12 + 14 + 15

offered-limited-capped-cir — Counter resource allocation:2

offered-limited-capped-cir is defined, the system creates four forwarding plane offered-output counters in the network processor and four discard counters in the traffic manager.

The **offered-limited-capped-cir** mode is similar to the **offered-profile-capped-cir** mode except that it combines **profile out** and **soft-out-of-profile** and eliminates the ‘offered-undefined’ statistic.

The impact of using **offered-limited-capped-cir** stat-mode while **profile-capped** mode is disabled are that one of the counting resources in the forwarding plane and traffic manager will not be used and **soft-in-profile** will be treated as ‘offered-in’ instead of ‘offered-undefined’.

The counters are used in the following manner:

1. ‘offered-in-that-stayed-green’ = profile in, soft-in-profile
2. ‘offered-in-that-turned-yellow-or-red’ = profile in, soft-in-profile
3. ‘offered-out-that-turned-green’ = soft-out-of-profile
4. ‘offered-out- that-turned-yellow-or-red’ = profile out, soft-out-of-profile
5. ‘dropped-in-that-stayed-green’ = Same as 1
6. ‘dropped-in-that-turned-yellow-or-red’ = Same as 2
7. ‘dropped-out-that-turned-green’ = Same as 3
8. ‘dropped-out-that-turned-yellow-or-red’ = Same as 4
9. ‘forwarded-in-that-stayed-green’ = Derived from 1 - 5
10. ‘forwarded-in-that-turned-yellow’ = Derived from 2 - 6
11. ‘forwarded-out-that-turned-green’ = Derived from 3 - 7

12. ‘forwarded-out-that-turned-yellow’ = Derived from 4 - 8

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

a. ‘offered-in’ = 1 + 2

b. ‘offered-out’ = 3 + 4

c. ‘discard-in’ = 5 + 7

d. ‘discard-out’ = 6 + 8

e. ‘forward-in’ = 9 + 11

f. ‘forward-out’ = 10 + 12

prec

Syntax **prec** *ip-prec-value* **fc** *fc-name* [**priority** {**high** | **low**}]
no prec *ip-prec-value*

Context config>qos>sap-ingress

Description This command explicitly sets the forwarding class or enqueuing priority when a packet is marked with an IP precedence value (*ip-prec-value)*. Adding an IP precedence rule on the policy forces packets that match the specified *ip-prec-value* to override the forwarding class and enqueuing priority based on the parameters included in the IP precedence rule.

When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy.

When the enqueuing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

The *ip-prec-value* is derived from the most significant three bits in the IP header ToS byte field (precedence bits). The three precedence bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop Quality-of-Service (QoS) behavior. The precedence bits are also part of the newer DiffServ Code Point (DSCP) method of mapping packets to QoS behavior. The DSCP uses the most significant six bits in the IP header ToS byte and so overlaps with the precedence bits. Both IP precedence and DSCP classification rules are supported. DSCP rules have a higher match priority than IP precedence rules and where a *dscp-name* DSCP value overlaps an *ip-prec-value*, the DSCP rule takes precedence.

The **no** form of the command removes the explicit IP precedence classification rule from the SAP ingress policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

Parameters *ip-prec-value* — The *ip-prec-value* is a required parameter that specifies the unique IP header ToS byte precedence bits value that will match the IP precedence rule. If the command is executed more than once with the same *ip-prec-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight IP precedence rules are allowed on a single policy.

The precedence is evaluated from the lowest to highest value.

Values 0 — 7

fc *fc-name* **—** The value given for the *fc-name* parameter must be one of the predefined forwarding classes in the system. Specifying the *fc-name* is optional. When a packet matches the rule the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The sub-class-name parameter is optional and used with the fc-name parameter to define a preexisting sub-class. The fc-name and sub-class-name parameters must be separated by a period (dot). If sub-class-name does not exist in the context of fc -name, an error will occur. If sub-class-name is removed using the no fc fc-name.sub-class-name force command, the default-fc command will automatically drop the sub-class-name and only use fc-name (the parent forwarding class for the sub-class) as the forwarding class.

Values fc: class[.sub-class]
class: be, l2, af, l1, h2, ef, h1, nc
sub-class: 29 characters max

Default Inherit (When **fc** is not defined, the rule preserves the previous forwarding class of the packet.)

priority — The priority parameter overrides the default enqueuing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule the enqueuing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueuing priority is inherited based on previous rule matches.

high — This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **high** for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

low — This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **low** for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Default Inherit (When priority is not defined, the rule preserves the previous enqueuing priority of the packet.)

Values high, low

prec

Syntax **prec** *ip-prec-value* **fc** *fc-name* [**hsmda-counter-override** *counter-id*] [**profile** {**in** | **out**}]
no prec *ip-prec-value*

Context config>qos>sap-egress

Description This command defines a specific IP precedence value that must be matched to perform the associated

reclassification actions. If an egress packet on the SAP matches the specified IP precedence value, the forwarding class, profile or HSMDA egress queue accounting behavior may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions. The default behavior for HSMDA queue accounting is to use the counters associated with the queue to which the packet is mapped.

The IP precedent bits used to match against prec reclassification rules come from the Type of Service (ToS) field within the IPv4 header. If the packet does not have an IPv4 header, prec based matching is not performed.

The reclassification actions from a prec reclassification rule may be overridden by a DHCP or IP flow matching events.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions. If a dhcp or ip-criteria match occurs after the prec match, the new forwarding class may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new fc, the fc from the prec match will be used.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior. If a dhcp or ip-criteria match occurs after the prec match, the new profile may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new profile, the profile from the prec match will be used.

The hsmda-counter-override keyword is optional. When specified and the egress SAP is created on an HSMDA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMDA uses each queues default queue counters for packets mapped to the queue. The hsmda-counter-override keyword is used to map the packet to an explicit exception counter. One of eight counters may be used. When the packet is mapped to an exception counter, the packet will not increment the queues discard or forwarding counters, instead the exception discard and forwarding counters will be used. The dscp based counter override decision may be overwritten by an ip-criteria reclassification rule match if the higher priority classification rule has an hsmda-counter-override action defined.

The **no** form of the command removes the reclassification rule from the SAP egress QoS policy.

Parameters **fc** *fc-name* **—** This keyword is optional. When specified, packets matching the IP precedence value will be explicitly reclassified to the forwarding class specified as fc-name regardless of the ingress classification decision. The explicit forwarding class reclassification may be overwritten by a higher priority dscp or ip-criteria reclassification match. The fc-name defined must be one of the eight forwarding classes supported by the system. To remove the forwarding class reclassification action for the specified prec-value, the prec command must be re-executed without the fc reclassification action defined.

Values be, l1, af, l2, h1, ef, h2 or nc

Default None

profile {**in** | **out}** **—** This keyword is optional. When specified, packets matching the IP precedence value will be explicitly reclassified to either in-profile or out-of-profile regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by a higher priority dscp or ip-criteria reclassification match. To remove the profile reclassification action for the specified prec-value, the prec command must be re-executed without the profile reclassification action defined.

in — The in parameter is mutually excusive to the out parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When in is specified, any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane. This value may be overwritten by an explicit profile action in a higher priority dscp or ip-criteria reclassification match.

out — The out parameter is mutually excusive to the in parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When out is specified, any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane. This value may be overwritten by an explicit profile action in a higher priority dscp or ip-criteria reclassification match.

hsmda-counter-override *counter-id* — The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the IP precedence value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The hsmda-counter-override action may be overwritten by an ip-criteria reclassification rule match. The specified counter-id must be specified as an integer between 1 and 8. To remove the ESMDA exception counter reclassification action for the specified DSCP value, the dscp command must be re-executed without the hsmda-counter-override reclassification action defined.

Values 1 —8

queue

Syntax **queue** *queue-id* [**multipoint**] [*queue-type*] [*queue-mode*] **pool** *pool-name*
queue *queue-id* [**multipoint**] [*queue-type*] **pool** *pool-name*
no queue *queue-id*

Context config>qos>sap-ingress
config>qos>sap-egress

Description This command creates the context to configure an ingress service access point (SAP) QoS policy queue.

Explicit definition of an ingress queue's hardware scheduler status is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or non-expedited nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (`nc`, `ef`, `h1` or `h2`), the queue is treated as an expedited queue by the hardware schedulers. When any non-expedited forwarding classes are mapped to the queue (`be`, `af`, `l1` or `l2`), the queue is treated as best effort (`be`) by the hardware schedulers. The expedited hardware schedulers are used to enforce expedited access to internal switch fabric destinations. The hardware status of the queue must be defined at the time of queue creation within the policy.

The **queue** command allows the creation of multipoint queues. Only multipoint queues can receive ingress packets that need flooding to multiple destinations. By separating the unicast for multipoint traffic at service ingress and handling the traffic on separate multipoint queues special handling of the multipoint traffic is possible. Each queue acts as an accounting and (optionally) shaping device offering precise control over potentially expensive multicast, broadcast and unknown unicast traffic. Only the back-end support of multipoint traffic (between the forwarding class and the queue based on forwarding type) needs to be defined. The individual classification rules used to place traffic into forwarding classes are not affected. Queues must be defined as multipoint at the time of creation within the policy.

The multipoint queues are for multipoint-destined service traffic. Within non-multipoint services, such as Epipe services, all traffic is considered unicast due to the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service will not be mapped to a multipoint service queue.

When an ingress SAP QoS policy with multipoint queues is applied to an Epipe SAP, the multipoint queues are not created. When an ingress SAP QoS policy with multipoint queues is applied to an IES SAP, a multipoint queue will be created when PIM is enabled on the IES interface.

Any billing or statistical queries about a multipoint queue on a non-multipoint service returns zero values. Any queue parameter information requested about a multipoint queue on a non-multipoint service returns the queue parameters in the policy. Buffers will not be allocated for multipoint queues on non-multipoint services. Buffer pool queries return zero values for actual buffers allocated and current buffer utilization.

The **no** form of this command removes the *queue-id* from the SAP ingress QoS policy and from any existing SAPs using the policy. If any forwarding class forwarding types are mapped to the queue, they revert to their default queues. When a queue is removed, any pending accounting information for each SAP queue created due to the definition of the queue in the policy is discarded.

The pool keyword is a create time parameter that allows the queue to receive buffers from an explicit buffer pool instead of the default buffer pool. The specified pool-name must have been explicitly created in a named-pool-policy and the policy must have been applied to the MDA or port on which the queue resides.

If the specified pool-name does not exist on the MDA, the queue will be treated as 'pool orphaned' and will be mapped to the appropriate default pool. Once the pool comes into existence on the MDA or port, the queue will be mapped to the new pool.

Once the queue is created within the policy, the pool command may be used to either remove the queue from the pool, or specify a new pool name association for the queue. The pool command does not appear in save or show command output. Instead, the current pool name for the queue will appear (or not appear) on the queue command output using the pool keyword.

Parameters *queue-id* — The *queue-id* for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

Values 1 — 32

queue-type — The **expedite**, **best-effort** and **auto-expedite** queue types are mutually exclusive to each other. Each defines the method that the system uses to service the queue from a hardware perspective. While parental virtual schedulers can be defined for the queue, they only enforce how the queue interacts for bandwidth with other queues associated with the same scheduler hierarchy. An internal mechanism that provides access rules when the queue is vying for bandwidth with queues in other virtual schedulers is also needed. A keyword must be specified at the time the queue is created in the SAP ingress policy. If an attempt to change the keyword after the queue is initially defined, an error is generated.

expedite — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

auto-expedite — This keyword allows the system to auto-define the way the queue is serviced by the hardware. When **auto-expedite** is defined on the queue, the queue is treated in an expedited manner when all forwarding classes mapped to the queue are configured as expedited types `nc`, `ef`, `h1` or `h2`.

When a single non-expedited forwarding class is mapped to the queue (`be`, `af`, `l1` and `l2`) the queue automatically falls back to non-expedited status.

Values expedite, best-effort, auto-expedite

Default auto-expedite

multipoint — This keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. If you attempt to map forwarding class unicast traffic to a multipoint queue, an error is generated and no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a pre-existing multipoint queue to edit *queue-id* parameters.

Values multipoint or not present

Default Present (the queue is created as non-multipoint)

queue-mode — Specifies the mode in which the queue is operating. This attribute is associated with the queue at the time of creation and cannot be modified thereafter.

Values **profile-mode**: When the queue is operating in the profile mode (or, the color aware mode), the queue tries to provide the appropriate bandwidth to the packets with different profiles. The profiles are assigned according to the configuration of the forwarding class or the sub-forwarding class.

priority-mode: The queue is capable of handling traffic differently with two distinct priorities. These priorities are assigned by the stages preceding the queueing framework in the system. In priority mode, the queue does not have the functionality to support the profiled traffic and in such cases the queue will have a degraded performance. However, the converse is not valid and a queue in profile mode should be capable of supporting the different priorities of traffic.

Default **priority-mode**

pool-name — The specified pool-name identifies a named pool where the policy will be applied. Each queue created within the system is tied to a physical port. When the policy is applied and the queue is created, the system will scan the named pools associated with the port to find the specified pool name. If the pool is not found on the port, the system will then look at named pools defined at the ports MDA level. If the pool name is not found on either the port or MDA, the queue will be marked as 'pool-orphaned' and will be mapped to the appropriate default pool. If the pool comes into existence, the queue will be moved from the default pool to the new named pool and the 'pool-orphaned' state will be cleared. The specified name must be an ASCII name string up to 16 characters long.

Values Any valid ASCII name string

Default None

The queue's pool association may only be removed by either re-executing the queue command without the pool keyword or by executing the no pool command within the queue's CLI context. When the pool name is removed, the queue will be placed on the appropriate default pool.

Service Ingress QoS Policy Forwarding Class Commands

broadcast-queue

Syntax **broadcast-queue** *queue-id* [**group** *queue-group-name*]

Context config>qos>sap-ingress>fc *fc-name*

Description This command overrides the default broadcast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all broadcast traffic on a SAP using this policy will be forwarded using the *queue-id*.

The broadcast forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of the command sets the broadcast forwarding type *queue-id* back to the default of tracking the multicast forwarding type queue mapping.

Parameters *queue-id* — The *queue-id* parameter must be an existing, multipoint queue defined in the `config>qos>sap-ingress` context.

Values Any valid multipoint queue ID in the policy including 2 through 32.

Default 11

group *queue-group-name* **—** This optional parameter is used to redirect the forwarding type within the forwarding class to the specified queue-id within the queue-group-name. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The *queue-group-name* are configured in the **config>qos>queue-group-templates** egress and ingress contexts.

de-1-out-profile

Syntax [**no**] **de-1-out-profile**

Context config>qos>sap-ingress>fc

Description This command, when enabled on a parent forwarding class, applies a color profile mode to the packets stored in the queue associated with this forwarding class. The queue associated with the parent forwarding class MUST be of type **profile-mode**.

When this QoS policy is applied to the ingress of a Frame Relay VLL SAP, the system will treat the received FR frames with DE bit set as out-of-profile regardless of their previous marking as the result of the default classification or on a match with an IP filter. It also adjusts the CIR of the ingress SAP queue to take into account out-of-profile frames which were sent while the SAP queue was in the "< CIR" state of the bucket. This makes sure that the CIR of the SAP is achieved in the long run.

All received DE=0 frames which are classified into this parent forwarding class or any of its sub-classes have their profile unchanged by enabling this option. That is the DE=0 frame profile could be undetermined

(default), in-profile, or out-of-profile as per previous classification. The DE=0 frames which have a profile of undetermined will be evaluated by the system CIR marking algorithm and will be marked appropriately.

The **priority** option if used has no effect. All FR VLL DE=1 frames have automatically their priority set to low while DE=0 frames have their priority set to high. Furthermore, DE=1 frames have drop-preference bit set in the internal header. The internal settings of the priority bit and of the drop-preference bit of the frame is independent of the use or not of the profile mode.

All other capabilities of the Fpipe service are maintained. This includes remarking of the DE bit on egress SAP, and FR PW control word on egress network port for the packets which were classified into "out-of-profile" at ingress SAP.

This **de-1-out-profile** keyword has an effect when applied to the ingress of a SAP which is part of an fpipe service. It can also be used on the ingress of an epipe or vpls SAP.

The **no** form of the command disables the color profile mode of operation on all SAPs this ingress QoS policy is applied.

Default no de-1-out-profile

egress-fc

Syntax **egress-fc** *fc-name*
no egress-fc

Context config>qos>sap-ingress>fc

Description This command configures the forwarding class to be used by the egress QOS processing. It overrides the forwarding class determined by ingress classification but no0t the QOS Policy Propagation via BGP.

The forwarding class and/or forwarding sub-class can be overriden.

The new egress forwarding class is applicable to both SAP egress and network egress.

Default no egress-fc

Parameters *fc-name* — Specifies the forwarding class name to be used by the egress QOS processing.

Default None. The fc name must be specified.

Values be, l2, af, l1, h2, ef, h1, nc

in-remark

Syntax **in-remark dscp** *dscp-name*
in-remark prec *ip-prec-value*
no in-remark

Context config>qos>sap-ingress>fc *fc-name*

Description This command is used in a SAP ingress QoS policy to define an explicit in-profile remark action for a forwarding class or sub-class. While the SAP ingress QoS policy may be applied to any SAP, the remarking functions are only enforced when the SAP is associated with an IP interface (IES or VPRN SAPs). When the

policy is applied to a Layer 2 SAP (i.e., Epipe or VPLS), the remarking definitions are silently ignored.

In the case where the policy is applied to a Layer 3 SAP, the in-profile remarking definition will be applied to packets that have been classified to the forwarding class or sub-class. It is possible for a packet to match a classification command that maps the packet to a particular forwarding class or sub-class only to have a more explicit (higher priority match) override the association. Only the highest priority match forwarding class or sub-class association will drive the in-profile marking.

The in-remark command is only applicable to ingress IP routed packets that are considered in-profile. The profile of a SAP ingress packet is affected by either the explicit in-profile/out-of-profile definitions or the ingress policing function applied to the packet. The following table shows the effect of the in-remark command on received SAP ingress packets. Within the in-profile IP packet's ToS field, either the six DSCP bits or the three precedence bits are remarked.

SAP Ingress Packet State	'in-remark' Command Effect
Non-Routed, Policed In-Profile	No Effect (non-routed packet)
Non-Routed, Policed Out-of-Profile	No Effect (non-routed packet)
Non-Routed, Explicit In-Profile	No Effect (non-routed packet)
Non-Routed, Explicit Out-of-Profile	No Effect (non-routed packet)
IP Routed, Policed In-Profile	in-remark value applied to IP header ToS field
IP Routed, Policed Out-of-Profile	No Effect (out-of-profile packet)
IP Routed, Explicit In-Profile	in-remark value applied to IP header ToS field
IP Routed, Explicit Out-of-Profile	No Effect (out-of-profile packet)

The **no** form of the command disables ingress remarking of in-profile packets classified to the forwarding class or sub-class.

Parameters **dscp** *dscp-name* **—** This parameter is one of two mutually exclusive settings that are applicable to the in-remark command. The in-remark command can be configured to either remark the DiffServ Code Point (DSCP) six bit value or the three Precedence bits. The dscp parameter specifies that the matching packets DSCP bits should be overridden with the value represented by dscp-name.

32 characters, maximum, The name specified by dscp-name is used to refer to the six bit value represented by dscp-name. It must be one of the predefined DSCP names defined on the system.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

Default None (an explicit valid DSCP name must be specified)

prec *ip-prec-value* **—** This parameter is one of two mutually exclusive settings that are applicable to the in-remark command. The in-remark command can be configured to either remark the DiffServ Code Point (DSCP) six bit value or the three Precedence bits. The prec parameter specifies that the matching packets Precedence bits should be overridden with the value represented by prec-value.

Values 0 — 7

Default None (an explicit precedence value must be specified)

multicast-queue

Syntax **multicast-queue** *queue-id* [**group** *queue-group-name*]

Context config>qos>sap-ingress>fc *fc-name*

Description This command overrides the default multicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all multicast traffic on a SAP using this policy is forwarded using the *queue-id*.

The multicast forwarding type includes the **unknown** unicast forwarding type and the **broadcast** forwarding type unless each is explicitly defined to a different multipoint queue. When the unknown and broadcast forwarding types are left as default, they will track the defined queue for the multicast forwarding type.

The **no** form of the command sets the multicast forwarding type *queue-id* back to the default queue for the forwarding class. If the **broadcast** and **unknown** forwarding types were not explicitly defined to a multipoint queue, they will also be set back to the default multipoint queue (queue 11).

Parameters *queue-id* — The *queue-id* parameter specified must be an existing, multipoint queue defined in the `config>qos>sap-ingress` context.

Values Any valid multipoint queue-ID in the policy including 2 through 32.

Default 11

group *queue-group-name* **—** This optional parameter is used to redirect the forwarding type within the forwarding class to the specified queue-id within the queue-group-name. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The *queue-group-name* are configured in the **config>qos>queue-group-templates** egress and ingress contexts.

out-remark

Syntax **out-remark dscp** *dscp-name*
out-remark prec *ip-prec-value*
no out-remark

Context config>qos>sap-ingress>fc *fc-name*

Description This command is used in a SAP ingress QoS policy to define an explicit out-of-profile remark action for a forwarding class or sub-class. While the SAP ingress QoS policy may be applied to any SAP, the remarking

functions are only enforced when the SAP is associated with an IP interface (IES or VPRN SAPs). When the policy is applied to a Layer 2 SAP (for example, Epipe or VPLS), the remarking definitions are silently ignored.

In the case where the policy is applied to a Layer 3 SAP, the out-of-profile remarking definition will be applied to packets that have been classified to the forwarding class or sub-class. It is possible for a packet to match a classification command that maps the packet to a particular forwarding class or sub-class only to have a more explicit (higher priority match) override the association. Only the highest priority match forwarding class or sub-class association will drive the out-of-profile marking.

The out-remark command is only applicable to ingress IP routed packets that are considered out-of-profile. The profile of a SAP ingress packet is affected by either the explicit in-profile/out-of-profile definitions or the ingress policing function applied to the packet. The following table shows the effect of the out-remark command on received SAP ingress packets. Within the out-of-profile IP packet's ToS field, either the six DSCP bits or the three precedence bits are remarked.

Table 35: Out-remark command effect

SAP Ingress Packet State	'out-remark' Command Effect
Non-Routed, Policed In-Profile	No Effect (non-routed packet)
Non-Routed, Policed Out-of-Profile	No Effect (non-routed packet)
Non-Routed, Explicit In-Profile	No Effect (non-routed packet)
Non-Routed, Explicit Out-of-Profile	No Effect (non-routed packet)
IP Routed, Policed In-Profile	No Effect (in-profile packet)
IP Routed, Policed Out-of-Profile	out-remark value applied to IP header ToS field
IP Routed, Explicit In-Profile	No Effect (in-of-profile packet)
IP Routed, Explicit Out-of-Profile	out-remark value applied to IP header ToS field

A packet that is explicitly remarked at ingress will not be affected by any egress remarking decision. Explicit ingress remarking has highest priority.

The **no** form of the command disables ingress remarking of out-of-profile packets classified to the forwarding class or sub-class.

Default none

Parameters **dscp** *dscp-name* **—** This parameter is one of two mutually exclusive settings that are applicable to the out-remark command. The out-remark command can be configured to either remark the DiffServ Code Point (DSCP) six bit value or the three Precedence bits. The dscp parameter specifies that the matching packets DSCP bits should be overridden with the value represented by dscp-name.

32 characters, maximumThe name specified by dscp-name is used to refer to the six bit value represented by dscp-name. It must be one of the predefined DSCP names defined on the system.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, e f, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57 , cp58, cp59, cp60, cp61, cp62, cp63

Default None (an explicit valid DSCP name must be specified)

prec *ip-prec-value* **—** This parameter is one of two mutually exclusive settings that are applicable to the out-remark command. The out-remark command can be configured to either remark the DiffServ Code Point (DSCP) six bit value or the three Precedence bits. The prec parameter specifies that the matching packets Precedence bits should be overridden with the value represented by prec-value.

The value specified by prec-value is used to overwrite the Precedence bits within a matching routed packets IP header ToS field.

Values 0 — 7

Default None (an explicit Precedence value must be specified)

An explicit dscp name or prec value must be specified for out-of-profile remarking to be applied.

policer

Syntax **policer** *policer-id* [**fp-redirect-group**]
no policer *policer-id*

multicast-policer

Syntax **multicast-policer** *policer-id* [**fp-redirect-group**]
no multicast-policer *policer-id*

broadcast-policer

Syntax **broasdcast-policer** *policer-id* [**fp-redirect-group**]
no broadcast-policer *policer-id*

unknown-policer

Syntax **unknown-policer** *policer-id* [**fp-redirect-group**]
no unknown-policer *policer-id*

profile

Syntax **profile** {**in** | **out**}
no profile

Context config>qos>sap-igress>fc

Description This command places a forwarding class or sub-class into a color aware profile mode. Normally, packets associated with a class are considered in-profile or out-of-profile solely based on the dynamic rate of the ingress queue relative to its CIR. Explicitly defining a class as in-profile or out-of-profile overrides this function by handling each packet with the defined profile state.

The profile command may only be executed when the forwarding class or the parent forwarding class (for a sub-class) is mapped to a queue that has been enabled to support color aware profile packets. The queue may

only be configured for profile-mode at the time the queue is created in the SAP ingress QoS policy.

A queue operating in profile-mode may support in-profile, out-of-profile and non-profiled packets simultaneously. However, the high and low priority classification actions are ignored when the queue is in profile-mode.

The **no** form of the command removes an explicit in-profile or out-of-profile configuration on a forwarding class or sub-class.

Default **no profile** — The default profile state of a forwarding class or sub-class is not to treat ingress packets as color aware. An explicit definition for in-profile or out-of-profile must be specified on the forwarding class or sub-class.

Parameters **in —** The **in** keyword is mutually exclusive to the **out** keyword. When the profile in command is executed, all packets associated with the class will be handled as in-profile. Packets explicitly handled as in-profile or out-of-profile still flow through the ingress service queue associated with the class to preserve order within flows. In-profile packets will count against the CIR of the queue, diminishing the amount of CIR available to other classes using the queue that are not configured with an explicit profile.

out — The **out** keyword is mutually exclusive to the **in** keyword. When the profile out command is executed, all packets associated with the class will be handled as out-of-profile. Packets explicitly handled as in-profile or out-of-profile still flow through the ingress service queue associated with the class to preserve order within flows. Out-of-profile packets will not count against the CIR of the queue, allowing other classes using the queue that are not configured with an explicit profile to be measured against the full CIR.

unknown-queue

Syntax **unknown-queue** *queue-id* [**group** *queue-group-name*]
no unknown-queue

Context config>qos>sap-ingress>fc *fc-name*

Description This command overrides the default unknown unicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the *queue-id*.

The unknown forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of this command sets the unknown forwarding type *queue-id* back to the default of tracking the multicast forwarding type queue mapping.

Parameters *queue-id* — Specifiesan existing multipoint queue defined in the **config>qos>sap-ingress** context.

Values Any valid multipoint *queue-id* in the policy including 2 through 32.

Default 11

group *queue-group-name* **—** This optional parameter is used to redirect the forwarding type within the forwarding class to the specified queue-id within the queue-group-name. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the

specified queue group. The *queue-group-name* are configured in the **config>qos>queue-group-templates** egress and ingress contexts.

queue

Syntax **queue** *queue-id* [{**group** *queue-group-name* [**instance** *instance-id*]} | **port-redirect-group-queue**]
no queue

Context config>qos>sap-ingress>fc
config>qos>sap-egress>fc

Description This command overrides the default queue mapping for **fc** fc-name. The specified queue-id must exist within the policy before the mapping can be made. Once the forwarding class mapping is executed, all traffic classified to fc-name on a SAP using this policy.

The **no** form of this command sets the queue-id back to the default queue for the forwarding class (queue 1).

Default **no queue**

Parameters *queue-id* — Specifies the SAP egress queue-id to be associated with the forwarding class. The queue-id must be an existing queue defined in sap-egress policy-id.

Values 1 — 8.

Default 1

group *queue-group-name* **—** This optional parameter is used to redirect the forwarding type within the forwarding class to the specified *queue-id* within the *queue-group-name*. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The queue-group-name are configured in the *config>qos>queue-group-templates* egress and ingress contexts. This parameter is used when policy based queue group redirection is desired. That is, the specific queue group to redirect to is named in the QoS policy.

instance *instance-id* **—** This parameter is used to specify the specific instance of a queue group with template queue-group-name to which this queue should be redirected. This parameter is only valid for queue groups on egress ports where policy based redirection is required.

Values 1 — 40960

Default 1

port-redirect-group-queue — This keyword is used to mark a given forwarding class queue for redirection to an egress queue group queue. This is only used when the specific queue group instance is assigned at the time the QOS policy is applied to the SAP. This redirection model is known as SAP based redirection.

hsmda-queues

Syntax **hsmda-queues**

Context config>qos>sap-egress

Description This command enables the context to configure queue definitions for use on SAPs or subscribers on HSMDAs. A single QoS policy simultaneously defines queues for both standard MDA and for HSMDA subscribers and SAPs. This allows the policy association decision to be ignorant of the type of hardware the SAP or subscriber is traversing.

queue

Syntax **queue** *queue-id* [**port-redirect-group-queue**]
no queue *queue-id*

Context config>qos>sap-egress>hsmda-queues

Description This command, within the QoS policy HSMDA-queues context, is a container for the configuration parameters controlling the behavior of an HSMDA queue. Unlike the standard QoS policy queue command, this command is not used to actually create or dynamically assign the queue to the object which the policy is applied. The queue identified by queue-id always exists on the SAP or subscriber context whether the command is executed or not. In the case of HSMDA SAPs and subscribers, all eight queues exist at the moment the sys- tem allocates an HSMDA queue group to the object (both ingress and egress).

Best-Effort, Expedited and Auto-Expedite Queue Behavior Based on Queue-ID

With standard service queues, the scheduling behavior relative to other queues is based on two items, the queues Best-Effort or Expedited nature and the dynamic rate of the queue relative to the defined CIR. HSMDA queues are handled differently. The create time auto-expedite and explicit expedite and best-effort qualifiers have been eliminated and instead the scheduling behavior is based solely on the queues identifier. Queues with a queue-id equal to 1 are placed in scheduling class 1. Queues with queue-id 2 are placed in scheduling class 2. And so on up to scheduling class 8. Each scheduling class is either mapped directly to a strict scheduling priority level based on the class ID, or the class may be placed into a weighted scheduling class group providing byte fair weighted round robin scheduling between the members of the group. Two weighted groups are supported and each may contain up to three consecutive scheduling classes. The weighed group assumes its highest member class's inherent strict scheduling level for scheduling purposes. Strict priority level 8 has the highest priority while strict level 1 has the lowest. When grouping of scheduling classes is defined, some of the strict levels will not be in use.

Single Type of HSMDA Queues

Another difference between HSMDA queues and standard service queues is the lack of Multipoint queues. At ingress, an HSMDA SAP or subscriber does not require multi-point queues since all forwarding types (broadcast, multicast, unicast and unknown) forward to a single destination, the ingress forwarding plane on the IOM. Instead of a possible eight queues per forwarding type (for a total of up to 32) within the SAP ingress QoS policy, the hsmda-queues node supports a maximum of eight queues.

Every HSMDA Queue Supports Profile Mode Implicitly

Unlike standard service queues, the HSMDA queues do not need to be placed into the special mode profile at create time in order to support ingress color aware policing. Each queue may handle in-profile, out-of-profile and profile undefined packets simultaneously. As with standard queues, the explicit profile of a packet is dependant on ingress sub-forwarding class to which the packet is mapped.

Queue sharing and redirection

Redirection to an egress port queue group specified for the HSMDA is possible using the port-redirect-group parameter. If this is specified, then packets are redirected to the queue-id in the HSMDA queue group instance named at the the time the egress QoS policy is applied to the SAP.

The **no** form of the command restores the defined queue-id to its default parameters. All HSMDA queues having the queue-id and associated with the QoS policy are re-initialized to default parameters.

Parameters *queue-id* — Defines the context of which of the eight ingress or egress queues will be entered for editing purposes.

port-redirect-group — This parameter is used to mark a given forwarding class queue for redirection to an egress port queue group. This is only used when the specific queue group instance is assigned at the time the qos policy is applied to the SAP. This redirection model is knowen as SAP based redirection.

packet-byte-offset

Syntax **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
no packet-byte-offset

Context config>qos>sap-egress>hsmda-queues

Description This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4 byte CRC (everything except the preamble and inter-frame gap). As an example, the packet-byte-offset command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter
- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter
- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based

and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 20 bytes may be added to the packet and up to 43 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command packet-byte-offset subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As inferred above, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. When the queue group represents the last-mile bandwidth constraints for a subscriber, the offset allows the HSMDA queue group to provide an accurate accounting to prevent overrun and underrun conditions for the subscriber. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscribers packets on an Ethernet aggregation network.

The packet-byte-offset value may be overridden for the HSMDA queue at the SAP or subscriber profile level.

The **no** form of the command removes any accounting size changes to packets handled by the queue. The command does not affect overrides that may exist on SAPs or subscriber profiles associated with the queue.

Parameters **add** *add-bytes* **—** Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The **add** keyword is mutually exclusive with the **subtract** keyword.

Values 1 — 31

subtract *sub-bytes* **—** Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The **subtract** keyword is mutually exclusive with the **add** keyword. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 1 — 32

wrr-policy

Syntax **wrr-policy** *wrr-policy-name*
no wrr-policy

Context config>qos>sap-egress>hsmda-queues

Description This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy to the HSMDA queue.

Parameters *wrr-policy-name* — Specifies the existing HSMDA WRR policy name to associate to the queue.

low-burst-max-class

Syntax **low-burst-max-class** *class-id*
no low-burst-max-class

Context config>qos>sap-egress>hsmda-queues>queue

Description This command assigns the low burst maximum class to associate with the HSMDA queue.

The **no** form of the command returns the class id for the queue to the default value.

Parameters *class-id* — Specifies the class identifier of the low burst max class for the HSMDA queue.

Values 1— 32

wrr-weight

Syntax **wrr-weight** *value*
no wrr-weight

Context config>qos>sap-egress>hsmda-queues>queue

Description This command assigns the weight value to the HSMDA queue.

The **no** form of the command returns the weight value for the queue to the default value.

Parameters *percentage* — Specifies the weight for the HSMDA queue.

Values 1— 32

slope-policy

Syntax **slope-policy** *hsmda-slope-policy-name*
no slope-policy

Context config>qos>sap-egress>hsmda-queues>queue

Description This command associates an existing HSMDA slope policy to the QoS policy HSMDA queue. The specified hsmda-slope-policy-name must exist for the command to succeed. If the policy name does not exist, the command has no effect on the existing slope policy association. Once a slope policy is associated with a QoS policy queue, subscriber profile override or SAP override, the slope policy cannot be removed from the system. Any edits to an associated slope policy are immediately applied to the queues using the slope policy.

Within the ingress and egress QoS policies, packets are classified as high priority or low-priority. For color aware policies, packets are also potentially classified as in-profile, out-of-profile or profile-undefined. Based on these classifications, packets are mapped to the RED slopes in the following manner:

Ingress Slope Mapping

- In-Profile — High Slope (priority ignored)

- Profile-Undefined, High Priority — High Slope
- Out-of-Profile Low Slope (priority ignored)
- Profile-Undefined, Low Priority — Low Slope

Egress Slope Mapping

- In-Profile from ingress — High Slope
- Out-of-Profile from ingress — Low Slope

The specified policy contains a value that defines the queue's MBS value (queue-mbs). This is the maximum depth of the queue specified in bytes where all packets start to discard. The high and low priority RED slopes provide congestion control mechanisms that react to the current depth of the queue and start a random discard that increases in probability as the queue depth increases. The start point and end point for each discard probability slope is defined as follows:

- Start-Utilization — This is defined as percentage of MBS and specifies where the discard probability for the slope begins to rise above 0%. (A corresponding Start-Probability parameter is not needed as the start probability is always 0%.
- Maximum-Utilization — This is also defined as a percentage of MBS and specifies where (based on MBS utilized) the discard probability rises to 100%. This is the first portion of the knee coordinates and is meaningless without the Maximum-Probability parameter.
- Maximum-Probability — This is defined as a percentage of discard probability and in conjunction with maximum-utilization completes the knee coordinate where the discard probability deviates from the slope and rises to 100%.

Up to 1024 HSMDA slope policies may be configured on a system.

The system maintains a slope policy named **hsmda-default** which acts as a default policy when an explicit slope policy has not been defined for an HSMDA queue. The default policy may be edited, but it cannot be deleted. If a no slope-policy hsmda-default command is executed, the default slope policy returns to the factory default settings. The factory default settings are as follows:

High Slope:

- Start-Utilization 100%
- Max-Utilization 100%
- Max-Probability 100%
- Shutdown

Low Slope:

- Start-Utilization 90%
- Max-Utilization 90%
- Max-Probability 1
- No Shutdown

Time-Average-Factor: 0

The **no** form of the command restores the association between the queue and the HSMDA default slope policy. The command has no immediate effect for queues that have a local override defined for the slope policy.

Parameters *hsmda-slope-policy-name* — Specifies an existing slope policy within the system. If a slope policy with the specified name does not exist, the slope-policy command will fail without modifying the slope behavior on the queue. Once a slope policy is associated with an HSMDA queue, the policy cannot be deleted.

Default hsmda-default

Service Ingress QoS Policy Entry Commands

action

Syntax **action** [**fc** *fc-name*] [**priority** {**high** | **low**}][**policer** *policer-id*]
no action

Context config>qos>sap-ingress>ip-criteria>entry
config>qos>sap-ingress>ipv6-criteria>entry
config>qos>sap-ingress>mac-criteria>entry

Description This mandatory command associates the forwarding class or enqueuing priority with specific IP, IPv6 or MAC criteria entry ID. The action command supports setting the forwarding class parameter to a sub-class. Packets that meet all match criteria within the entry have their forwarding class and enqueuing priority overridden based on the parameters included in the **action** parameters. When the forwarding class is not specified in the **action** command syntax, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueuing priority is not specified in the action, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

When a policer is specified in the action, a matching packet is directed to the configured policer instead of the policer/queue assigned to the forwarding class of the packet.

The **action** command must be executed for the match criteria to be added to the active list of entries. If the entry is designed to prevent more explicit (higher entry ID) entries from matching certain packets, the **fc** *fc-name* and **match** *protocol* fields should not be defined when executing action. This allows packets matching the entry to preserve the forwarding class and enqueuing priority derived from previous classification rules.

Each time action is executed on a specific entry ID, the previous entered values for **fc** *fc-name* and **priority** areoverridden with the newly defined parameters or inherits previous matches when a parameter is omitted.

The **no** form of the command removes the entry from the active entry list. Removing an entry on a policy immediately removes the entry from all SAPs using the policy. All previous parameters for the action is lost.

Default Action specified by the **default-fc**.

Parameters **fc** *fc-name* — The value given for **fc** *fc-name* must be one of the predefined forwarding classes in the system. Specifying the **fc** *fc-name* is required. When a packet matches the rule, the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The sub-class-name parameter is optional and used with the fc-name parameter to define a preexisting sub-class. The fc-name and sub-class-name parameters must be separated by a period (dot). If sub-class-name does not exist in the context of fc -name, an error will occur. If sub-class-name is removed using the no fc fc-name.sub-class-name force command, the default-fc command will automatically drop the sub-class-name and only use fc-name (the parent forwarding class for the sub-class) as the forwarding class.

Values fc: class[.sub-class]
class: be, l2, af, l1, h2, ef, h1, nc
sub-class: 29 characters max

Default Inherit (When **fc** *fc-name* is not defined, the rule preserves the previous forwarding class of the packet.)

priority — The **priority** parameter overrides the default enqueuing priority for all packets received on a SAP using this policy that match this rule. Specifying the priority (**high** or **low**) is optional. When a packet matches the rule, the enqueuing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueuing priority is inherited based on previous rule matches.

Default Inherit (When the **priority** (**high** or **low**) is not defined, the rule preserves the previous enqueuing priority of the packet)

high — The **high** parameter is used in conjunction with the **priority** parameter. Setting the **priority** enqueuing parameter to **high** for a packet increases the likelihood to enqueue the packet when the queue is congested. The enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the queue, the significance of the enqueuing priority is lost.

low — The **low** parameter is used in conjunction with the **priority** parameter. Setting the **priority** enqueuing parameter to **low** for a packet decreases the likelihood to enqueue the packet when the queue is congested. The enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Default Inherit

policer-id — A valid policer-id must be specified. The parameter policer-id references a policer-id that has already been created within the sap-ingress QoS policy.

Values 1—32

Default none

action

Syntax **action** [**fc** *fc-name*] [**hsmda-counter-override** *counter-id*] [**profile** {**in** | **out**}]
no action

Context config>qos>sap-egress>ip-criteria>entry
config>qos>sap-egress>ipv6-criteria>entry

Description This command defines the reclassification actions that should be performed on any packet matching the defined IP flow criteria within the entries match node. When defined under the ip-criteria context, the reclassification ony applies to IPv4 packets. When defined under the ipv6-criteria context, the reclassification only applies to IPv6 packets.

If an egress packet on the SAP matches the specified IP flow entry, the forwarding class, profile egress queue accounting behavior may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions. The default behavior for HSMDA queue accounting is to use the counters associated with the queue to which the packet is mapped. Matching an IP flow reclassification entry will override all IP precedence or DSCP based reclassification rule actions when

an explicit reclassification action is defined for the entry.

When an IP flow entry is first created, the entry will have no explicit behavior defined as the reclassification actions to be performed. In show and info commands, the entry will display no action as the specified reclassification action for the entry. When the entry is defined with no action, the entry will not be populated in the IP flow reclassification list used to evaluate packets egressing a SAP with the SAP egress policy defined. An IP flow reclassification entry is only added to the evaluation list when the action command for the entry is executed either with explicit reclassification entries or without any actions defined. Specifying action without any trailing reclassification actions allows packets matching the entry to exist the evaluation list without matching entries lower in the list. Executing no action on an entry removes the entry from the evaluation list and also removes any explicitly defined reclassification actions associated with the entry.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior.

The **hsmda-counter-override** keyword is optional. When specified and the egress SAP is created on an HSMDA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMDA uses each queues default queue counters for packets mapped to the queue. The hsmda-counter-override keyword is used to map the packet to an explicit exception counter. One of eight counters may be used. When the packet is mapped to an exception counter, the packet will not increment the queues discard or forwarding counters, instead the exception discard and forwarding counters will be used.

The **no** form of this command removes all reclassification actions from the IP flow reclassification entry and also removes the entry from the evaluation list. An entry removed from the evaluation list will not be matched to any packets egress a SAP associated with the SAP egress QoS policy.

Default Action specified by the **default-fc**.

Parameters **fc** *fc-name* **—** The fc reclassification action is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to the forwarding class specified as fc-name regardless of the ingress classification decision. The fc-name defined must be one of the eight forwarding classes supported by the system. To remove the forwarding class reclassification action for the IP flow entry, the action command must be re-executed without the fc reclassification action defined.

Values fc: class[.sub-class]
class: be, l2, af, l1, h2, ef, h1, nc
sub-class: 29 characters max

Default none

profile {**in** | **out**} **—** The profile reclassification action is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to either in-profile or out-of-profile regardless of the ingress profiling decision. To remove the profile reclassification action for the IP flow reclassification entry, the action command must be re-executed without the profile reclassification action defined.

in — The in parameter is mutually excusive to the out parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When **in** is specified,

any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

out — The out parameter is mutually excusive to the in parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When out is specified, any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

hsmda-counter-override *counter-id* **—** The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the IP precedence value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The hsmda-counter-override action may be overwritten by an ip-criteria reclassification rule match. The specified counter-id must be specified as an integer between 1 and 8. To remove the ESMDA exception counter reclassification action for the specified dscp-value, the dscp command must be re-executed without the hsmda-counter-override reclassification action defined.

Values 1—8

Default None

entry

Syntax **entry** *entry-id* [**create**]
no entry *entry-id*

Context config>qos>sap-ingress>ip-criteria
config>qos>sap-egress>ip-criteria
config>qos>sap-ingress>ipv6-criteria
config>qos>sap-ingress>mac-criteria

Description This command is used to create or edit an IP or IPv6 criteria entry for the policy. Multiple entries can be created using unique *entry-id* numbers.

The list of flow criteria is evaluated in a top down fashion with the lowest entry ID at the top and the highest entry ID at the bottom. If the defined match criteria for an entry within the list matches the information in the egress packet, the system stops matching the packet against the list and performs the matching entries reclassification actions. If none of the entries match the packet, the IP flow reclassification list has no effect on the packet.

An entry is not populated in the list unless the action command is executed for the entry. An entry that is not populated in the list has no effect on egress packets. If the action command is executed without any explicit reclassification actions specified, the entry is populated in the list allowing packets matching the entry to exit the list, preventing them from matching entries lower in the list. Since this is the only flow reclassification entry that the packet matched and this entry explicitly states that no reclassification action is to be performed, the matching packet will not be reclassified.

The **no** form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services where that policy is applied.

Default none

Parameters *entry-id* — The *entry-id*, expressed as an integer, uniquely identifies a match criterion and the corresponding action. It is recommended that multiple entries be given *entry-ids* in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword **action fc** *fc-name* for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive.

Default none

Values 1— 65535

create — Required parameter when creating a flow entry when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the flow entry already exists.

match

Syntax [**no**] **match** [**protocol** *protocol-id*]

Context config>qos>sap-egress>ip-criteria>entry
config>qos>sap-ingress>ip-criteria>entry

Description This command creates a context to configure match criteria for SAP QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, then all criteria must be satisfied (AND function) before the action associated with the match is executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

It is possible that a SAP policy includes the **dscp** map command, the **dot1p** map command, and an IP match criteria. When multiple matches occur for the traffic, the order of precedence is used to arrive at the final action. The order of precedence is as follows:

1. 802.1p bits
2. DSCP
3. IP Quintuple or MAC headers

The **no** form of this command removes the match criteria for the *entry-id*.

Parameters **protocol** *protocol-id* — Specifies an IP protocol to be used as a SAP QoS policy match criterion.

The protocol type such as TCP / UDP / OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17).

Values protocol-id: 0 — 255 protocol numbers accepted in DHB

keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp,

ip, ipv6, ipv6-frag, ipv6-icmp, ipv6-no-nxt, ipv6-opts, ipv6-route, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp* — udp/tcp wildcard

Table 36: IP Protocol Names

Protocol	Protocol ID	Description
icmp	1	Internet Control Message
igmp	2	Internet Group Management
ip	4	IP in IP (encapsulation)
tcp	6	Transmission Control
egp	8	Exterior Gateway Protocol
igp	9	any private interior gateway (used by Cisco for their IGRP)
udp	17	User Datagram
rdp	27	Reliable Data Protocol
ipv6	41	IPv6
ipv6-route	43	Routing Header for IPv6
ipv6-frag	44	Fragment Header for IPv6
idrp	45	Inter-Domain Routing Protocol
rsvp	46	Reservation Protocol
gre	47	General Routing Encapsulation
ipv6-icmp	58	ICMP for IPv6
ipv6-no-nxt	59	No Next Header for IPv6
ipv6-opts	60	Destination Options for IPv6
iso-ip	80	ISO Internet Protocol
eigrp	88	EIGRP
ospf-igp	89	OSPFIGP
ether-ip	97	Ethernet-within-IP Encapsulation
encap	98	Encapsulation Header
pnni	102	PNNI over IP
pim	103	Protocol Independent Multicast
vrrp	112	Virtual Router Redundancy Protocol
l2tp	115	Layer Two Tunneling Protocol

Table 36: IP Protocol Names (Continued)

Protocol	Protocol ID	Description
stp	118	Schedule Transfer Protocol
ptp	123	Performance Transparency Protocol
isis	124	ISIS over IPv4
crtp	126	Combat Radio Transport Protocol
crudp	127	Combat Radio User Datagram

match

Syntax **match** [**next-header** *next-header*]
no match

Context config>qos>sap-ingress>ipv6-criteria>entry
config>qos>sap-egress>ipv6-criteria>entry

Description This command creates a context to configure match criteria for ingress SAP QoS policy match IPv6 criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, then all criteria must be satisfied (AND function) before the action associated with the match is executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

It is possible that a SAP ingress policy includes the **dscp** map command, the **dot1p** map command, and an IPv6 match criteria. When multiple matches occur for the traffic, the order of precedence is used to arrive at the final action. The order of precedence is as follows:

1. 802.1p bits
2. DSCP
3. IP Quintuple or MAC headers

The **no** form of this command removes the match criteria for the *entry-id*.

Parameters **next-header** *next-header* — Specifies the next meader to match.

The protocol type such as TCP / UDP / OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17).

Values protocol numbers accepted in DHB: 0 — 42, 45 — 49, 52 — 59, 61 — 255
keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp
* — udp/tcp wildcard

match

Syntax **match** [**frame-type** {**802dot3 | 802dot2-llc | 802dot2-snap | ethernet-II | atm**}]
no match

Context config>qos>sap-ingress>mac-criteria>entry

Description This command creates a context for entering/editing match MAC criteria for ingress SAP QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured then all criteria must be satisfied (AND function) before the action associated with the match will be executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

The **no** form of the command removes the match criteria for the *entry-id*.

Parameters **frame-type** *keyword* **—** The **frame-type** keyword configures an Ethernet frame type or an ATM frame type to be used for the MAC filter match criteria.

Default 802dot3

Values 802dot3, 802dot2-llc, 802dot2-snap, ethernet_II, atm

802dot3 — Specifies the frame type is Ethernet IEEE 802.3.

802dot2-llc — Specifies the frame type is Ethernet IEEE 802.2 LLC.

802dot2-snap — Specifies the frame type is Ethernet IEEE 802.2 SNAP.

ethernet_II — Specifies the frame type is Ethernet Type II.

atm — Specifies the frame type as ATM cell. Note that the user is not allowed to configure entries with frame type of atm and a frame type of other supported values in the same QoS policy.

atm-vci

Syntax **atm-vci** *vci-value*
no atm-vci

Context config>qos>sap-ingress>mac-criteria>entry>match

Description This command configures a VCI based filter entry in the SAP ingress QoS policy.

This new criterion has only take affect when applied to a VPI SAP of an apipe VLL service of type atm-vpc. The application of this criterion to the ATM SAP of any other ATM VLL service, any other VLL service, VPLS service, or IES/VPRN service has no effect.

The user is not allowed to configure a MAC matching criterion other than atm-vci once a MAC criteria filter entry that includes the frame type of atm has been configured.

When the policy is applied to the ingress ATM VPI SAP of an atm-vpc VLL service and a received packet matches the VCI value configured in the atm-vci parameter, it is assigned the FC in the fc option of the action part of the filter. This determines which forwarding class queue this packet will be stored. Note that if

the user entered a priority value in the priority option, it is ignored as the priority and profile of ATM VLL service packets is solely determined based on the ATM conformance definition configured in the ATM QoS traffic descriptor profile applied to this ATM SAP.

On egress ATM SAP, the Q-chip will queue the packet on the egress SAP queue corresponding to the packet's FC and forwards the packet to the ATM MDA (CMA). The ATM MDA (CMA) stores the individual cells it in the VP queue corresponding to the SAP.

It is strongly recommended that the user does not enable cell-concatenation on the spoke-SDP when a VCI QoS filter is applied to the SAP. The filter will match against the VCI in the header of the first cell in the concatenated packet. Cell concatenation is disabled by default on a spoke-sdp of all ATM VLL service types.

The **no** form of this command removes the VCI value as the match criterion.

Parameters *vci-value* — The value of the VCI field in the received ATM cell header.

Values 1, 2, 5 — 65535

IP QoS Policy Match Commands

dscp

Syntax **dscp**
no dscp

Context config>qos>sap-ingress>ip-criteria>entry>match
config>qos>sap-egress>ip-criteria>entry>match
config>qos>sap-ingress>ipv6-criteria>entry>match
config>qos>sap-egress>ipv6-criteria>entry>match

Description This command configures a DiffServ Code Point (DSCP) code point to be used for remarking of packets from the specified FC. If the optional in/out-profile is specified, the command will remark different DSCP code points depending on whether the packet was classified to be in or out-of-profile ingress to the node.

The **no** form of this command removes the DSCP match criterion.

Default none

Parameters *dscp-name* — Specifies a dscp name that has been previously mapped to a value using the **dscp-name** command. The DiffServ code point can only be specified by its name.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

hsmda

Syntax **hsmda**

Context config>qos>sap-egress>fc

Description This command defines how packets matching the forwarding class will be mapped to an HSMDA queue ID. The SAP QoS policies simultaneously support both standard service queue mappings and ESDMA queue mappings for the same forwarding class and the hsmda node is used to separate the HSMDA mappings from the standard mappings This allows the same QoS policy to be used on a standard MDA attached SAP and an HSMDA attached SAP.

queue

Syntax **queue** [1..8]
no queue

Context config>qos>sap-egress>fc>hsmda

Description This command specifies the HSMDA queue mapping for all packets in point-to-point services and unicast destined packets in multipoint services. Point-to-point services include epipe and other VLL type services. Multipoint services include IES, VPLS and VPRN services. The queue command does not apply to multicast, broadcast or unknown unicast packets within multipoint services (the multicast, broadcast and unknown commands must be used to define the queue mapping for non-unicast packets within a forwarding class). For Epipe, the **queue** *queue-id* mapping applies to all packets, regardless of the packets destination MAC address.

Each forwarding class has a default queue ID based on the intrinsic hierarchy between the forwarding classes as represented in Table 37. Executing the queue command within the HSMDA context of a forwarding class with a different queue ID than the default overrides the default mapping. Multiple forwarding classes may be mapped to the same HSMDA queue ID.

Table 37: Default FC HSMDA Queue ID Mappings

Forwarding Class	Default HSMDA Queue ID
NC	queue 8
H1	queue 7
EF	queue 6
H2	queue 5
L1	queue 4
AF	queue 3
L2	queue 2
BE	queue 1

Table 38 presents the way that packets are mapped to queues based on the type of service and the various forwarding types.

Table 38: Ingress HSMDA Queue Mapping Behavior Based on Forwarding Type

		Queue Mappings For Each Forwarding Type		
Service Type	**Queue**	**Broadcast**	**Multicast**	**Unknown**
Epipe	All packets matching the FC	None	None	None

Table 38: Ingress HSMDA Queue Mapping Behavior Based on Forwarding Type

	Queue Mappings For Each Forwarding Type			
IES	All packets matching the FC	Packets with Broadcast DA	IP Multicast Packets	None
VPLS	All packets matching the FC	Packets with Broadcast DA	Packets with Multicast DA	Packets with Unicast DA but Unknown in FIB
VPRN	All packets matching the FC	Packets with Broadcast DA	IP Multicast Packets	None

The forwarding class queue mappings may be modified at anytime. The sub-forwarding classes inherit the parent forwarding classes queue mappings.

The no form of the command returns the HSMDA queue mapping for queue to the default mapping for the forwarding class.

Parameters *queue-id* — Configures a specific HSMDA queue.

Values 1 — 8

BE Default:	1
L2 Default:	2
AF Default:	3
L1 Default:	4
H2 Default:	5
EF Default:	6
H1 Default:	7
NC Default:	8

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
no mbs

Context config>qos>sap-egress>hsmda-queues>queues

Description This command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low priority violate threshold. For egress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold.

The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the

policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by **high-prio-only** is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.

The no form of this command returns the policer to its default MBS size.

Default None

Parameters *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

byte — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobyte — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Values 1—39321600

Default **kilobyte**

dst-ip

Syntax **dst-ip** {*ip-address/mask* | *ip-address netmask*}
no dst-ip

Context config>qos>sap-ingress>ip-criteria>entry>match
config>qos>sap-egress>ip-criteria>entry>match
config>qos>sap-ingress>ipv6-criteria>entry>match
config>qos>sap-egress>ipv6-criteria>entry>match

Description This command configures a destination address range to be used as a SAP QoS policy match criterion.

To match on the destination address, specify the address and its associated mask, e.g., 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

The **no** form of this command removes the destination IP address match criterion.

Default none

Parameters *ip-address* — The IP address of the destination IP or IPv6 interface. This address must be unique within the subnet and specified in dotted decimal notation.

Values

ip-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x	(eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 — FFFF]H
	d:	[0 — 255]D
prefix-length:	1 — 128	

dst-port

Syntax **dst-port** {**lt** | **gt** | **eq**} *dst-port-number*
dst-port range *start end*
no dst-port

Context config>qos>sap-ingress
config>qos>sap-ingress>ip-criteria>entry>match
config>qos>sap-egress>ip-criteria>entry>match
config>qos>sap-ingress>ipv6-criteria>entry>match
config>qos>sap-egress>ipv6-criteria>entry>match

Description This command configures a destination TCP or UDP port number or port range for a SAP QoS policy match criterion.

The **no** form of this command removes the destination port match criterion.

Default none

Parameters **lt** | **gt** | **eq** *dst-port-number* **—** The TCP or UDP port numbers to match specified as less than (**lt**), greater than (**gt**) or equal to (**eq**) to the destination port value specified as a decimal integer.

Values 1 — 65535 (decimal)

range *start end* **—** The range of TCP or UDP port values to match specified as between the *start* and *end* destination port values inclusive.

Values 1 — 65535 (decimal)

fragment

Syntax **fragment** {**true** | **false**}
no fragment

Context config>qos>sap-ingress>ip-criteria>entry>match
config>qos>sap-egress>ip-criteria>entry>match

Description This command configures fragmented or non-fragmented IP packets as a SAP QoS policy match criterion.

The **no** form of this command removes the match criterion and matches all packets regardless of whether they are fragmented or not.

Default no fragment

Parameters **true —** Configures a match on all fragmented IP packets. A match will occur for all packets that have either the MF (more fragment) bit set OR have the Fragment Offset field of the IP header set to a non-zero value.

false — Configures a match on all non-fragmented IP packets. Non-fragmented IP packets are packets that have the MF bit set to zero and have the Fragment Offset field also set to zero.

src-ip

Syntax **src-ip** *{ip-address/mask* | **ip-address** *ipv4-address-mask*|**ip-prefix-list** *prefix-list-name*]}
src-ip {*ipv6-address/prefix-length* | **ipv6-address** *ipv6-address-mask*}
no src-ip

Context config>qos>sap-ingress>ip-criteria>entry>match
config>qos>sap-egress>ip-criteria>entry>match
config>qos>sap-ingress>ipv6-criteria>entry>match
config>qos>sap-egress>ipv6-criteria>entry>match

Description This command configures a source IPv4 or IPv6 address range to be used as an SAP QoS policy match criterion.

To match on the source IPv4 or IPv6 address, specify the address and its associated mask, e.g. 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used for IPv4.

The **no** form of the command removes the source IPv4 or IPv6 address match criterion.

Default No source IP match criterion.

Parameters ip-address — Specifies the source IPv4 address specified in dotted decimal notation.

Values ip-address: a.b.c.d

mask — Specifies the length in bits of the subnet mask.

Values 1 — 32

ipv4-address-mask — Specifies the subnet mask in dotted decimal notation.

Values a.b.c.d (dotted quad equivalent of mask length)

ipv6-address — Specifies the IPv6 prefix for the IP match criterion in hex digits.

Values ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D

prefix — Specifies the IPv6 prefix length for the ipv6-address expressed as a decimal integer.

Values 1 — 128

mask — Specifies the eight 16-bit hexadecimal pieces representing bit match criteria.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

src-port

Syntax **src-port** {**lt** | **gt** | **eq**} *src-port-number*
src-port range *start end*
no src-port

Context config>qos>sap-ingress>ip-criteria>entry>match
config>qos>sap-egress>ip-criteria>entry>match
config>qos>sap-ingress>ipv6-criteria>entry>match

Description This command configures a source TCP or UDP port number or port range for a SAP QoS policy match criterion.

The **no** form of this command removes the source port match criterion.

Default No src-port match criterion.

Parameters **lt** | **gt** | **eq** *src-port-number* **—** The TCP or UDP port numbers to match specified as less than (**lt**), greater than (**gt**) or equal to (**eq**) to the source port value specified as a decimal integer.

Values 1 — 65535 (decimal)

range *start end* **—** The range of TCP or UDP port values to match specified as between the *start* and *end* source port values inclusive.

Values 1 — 65535 (decimal)

Service Ingress MAC QoS Policy Match Commands

dot1p

Syntax **dot1p** *dot1p-value* [*dot1p-mask*]
no dot1p

Context config>qos>sap-ingress>mac-criteria>entry

Description The IEEE 802.1p value to be used as the match criterion.

Use the **no** form of this command to remove the dot1p value as the match criterion.

Default None

Parameters *dot1p-value* — Enter the IEEE 802.1p value in decimal.

Values 0 — 7

dot1pmask — This 3-bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	D	4
Hexadecimal	0xH	0x4
Binary	0bBBB	0b100

To select a range from 4 up to 7 specify *p-value* of 4 and a *mask* of 0b100 for value and mask.

Default 7 (decimal) (exact match)

Values 1 — 7 (decimal)

dsap

Syntax **dsap** *dsap-value* [*dsap-mask*]
no dsap

Context config>qos>sap-ingress>mac-criteria>entry

Description Configures an Ethernet 802.2 LLC DSAP value or range for an ingress SAP QoS policy match criterion.

This is a one-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.

Use the no form of this command to remove the dsap value as the match criterion.

Default None

Parameters *dsap-value* — The 8-bit dsap match criteria value in hexadecimal.

Values 0x00 — 0xFF (hex)

dsap-mask — This is optional and can be used when specifying a range of dsap values to use as the match criteria.

This 8 bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	DDD	240
Hexadecimal	0xHH	0xF0
Binary	0bBBBBBBBB	0b11110000

Default FF (hex) (exact match)

Values 0x00 — 0xFF (hex)

dst-mac

Syntax **dst-mac** *ieee-address* [*ieee-address-mask*]
no dst-mac

Context config>qos>sap-ingress>mac-criteria>entry

Description Configures a destination MAC address or range to be used as a Service Ingress QoS policy match criterion.

The no form of this command removes the destination mac address as the match criterion.

Default none

Parameters *ieee-address* — The MAC address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask — A 48-bit mask to match a range of MAC address values.

This 48-bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHH	0xFFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

All packets with a source MAC OUI value of 00-03-FA subject to a match condition should be specified as: 0003FA000000 0x0FFFFF000000

Default 0xFFFFFFFFFFFF (hex) (exact match)

Values 0x00000000000000 — 0xFFFFFFFFFFFF (hex)

etype

Syntax **etype** *etype-value*
no etype

Context config>qos>sap-ingress>mac-criteria>entry

Description Configures an Ethernet type II value to be used as a service ingress QoS policy match criterion.

The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame. For e.g. 0800 is used to identify the IP v4 packets.

The Ethernet type field is used by the Ethernet version-II frames. IEEE 802.3 Ethernet frames do not use the type field. For IEEE 802.3 frames use the dsap, ssap or snap-pid fields as match criteria.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The no form of this command removes the previously entered etype field as the match criteria.

Default None

Parameters *etype-value* — The Ethernet type II frame Ethertype value to be used as a match criterion expressed in hexadecimal.

Values 0x0600 — 0xFFFF

snap-oui

Syntax **snap-oui** {**zero** | **non-zero**}
no snap-oui

Context config>qos>sap-ingress>mac-criteria>entry

Description Configures an IEEE 802.3 LLC SNAP Ethernet frame OUI zero or non-zero value to be used as a service ingress QoS policy match criterion.

The **no** form of this command removes the criterion from the match criteria.

Default none

Parameters **zero —** Specifies to match packets with the three-byte OUI field in the SNAP-ID set to zero.

non-zero — Specifies to match packets with the three-byte OUI field in the SNAP-ID not set to zero.

snap-pid

Syntax **snap-pid** *snap-pid*
no snap-pid

Context config>qos>sap-ingress>mac-criteria>entry

Description Configures an IEEE 802.3 LLC SNAP Ethernet frame PID value to be used as a service ingress QoS policy match criterion.

This is a two-byte protocol id that is part of the IEEE 802.3 LLC SNAP Ethernet Frame that follows the three-byte OUI field.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.

Note: **snap-pid** match criteria is independent of the OUI field within the SNAP header. Two packets with different three-byte OUI fields but the same PID field will both match the same policy entry based on a snap-pid match criteria.

The **no** form of this command removes the snap-pid value as the match criteria.

Default none

Parameters *smap-pid* — The two-byte snap-pid value to be used as a match criterion in hexadecimal.

Values 0x0000 — 0xFFFF

src-mac

Syntax **src-mac** *ieee-address* [*ieee-address-mask*]
no src-mac

Context config>qos>sap-ingress>mac-criteria>entry

Description This command configures a source MAC address or range to be used as a service ingress QoS policy match criterion.

The **no** form of this command removes the source mac as the match criteria.

Default none

Parameters *ieee-address* — Enter the 48-bit IEEE mac address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask — This 48-bit mask can be configured using:

This 48 bit mask can be configured using the following formats

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDD	281474959933440

Format Style	Format Syntax	Example
Hexadecimal	0xHHHHHHHHHHHH	0x0FFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

To configure all packets with a source MAC OUI value of 00-03-FA are subject to a match condition, then the entry should be specified as: 003FA000000 0xFFFFFF000000

Default 0xFFFFFFFFFFFF (hex) (exact match)

Values 0x00000000000000 — 0xFFFFFFFFFFFF (hex)

ssap

Syntax **ssap** *ssap-value* [*ssap-mask*]
no ssap

Context config>qos>sap-ingress>mac-criteria>entry

Description This command configures an Ethernet 802.2 LLC SSAP value or range for an ingress SAP QoS policy match criterion.

This is a one-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The **no** form of this command removes the ssap match criterion.

Default none

Parameters *ssap-value* — The 8-bit ssap match criteria value in hex.

Values 0x00 — 0xFF (hex)

ssap-mask — This is optional and can be used when specifying a range of ssap values to use as the match criteria.

This 8 bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	DDD	240
Hexadecimal	0xHH	0xF0
Binary	0bBBBBBBBB	0b11110000

Default none

Values 0x00 — 0xFF

Service Egress QoS Policy Forwarding Class Commands

fc

Syntax **fc** *fc-name*
no fc *fc-name*

Context config>qos>sap-egress

Description The **fc** fc-*name* node within the SAP egress QoS policy is used to contain the explicitly defined queue mapping and dot1p marking commands for *fc-name*. When the mapping for *fc-name* points to the default queue and the dot1p marking is not defined, the node for *fc-name* is not displayed in the **show configuration** or **save configuration** output unless the detail option is specified.

The **no** form of the command removes the explicit queue mapping and dot1p marking commands for *fc-name*. The queue mapping reverts to the default queue for *fc-name* and the dot1p marking (if appropriate) uses the default of 0.

Default none

Parameters *fc-name* — This parameter specifies the forwarding class queue mapping or dot1p marking is to be edited. The value given for *fc-name* must be one of the predefined forwarding classes in the system.

Values `be`, `l2`, `af`, `l1`, `h2`, `ef`, `h1`, `nc`

parent-location

Syntax **parent-location** {**default**|**sla**}
no parent-location

Context config>qos>sap-egress

Description This command determines the expected location of the parent schedulers for queues configured with a parent command within the sap-egress policy. All parent schedulers must be configured within a scheduler-policy applied at the location corresponding to the parent-location parameter.

If a parent scheduler name does not exist at the specified location, the queue will not be parented and will be orphaned.

The **no** form of the command reverts to the default.

Default default

Parameters **default —** When the sap-egress policy is applied to an sla-profile for a subscriber, the parent schedulers of the queues need to be configured in the scheduler-policy applied to the subscriber's sub-profile.

When the sap-egress policy is applied to a SAP, the parent schedulers of the queues need to be configured in the scheduler-policy applied to the SAP or the multi-service site.

sla — When the sap-egress policy is applied to an sla-profile for a subscriber, the parent schedulers of the queues need to be configured in the scheduler-policy applied to the same sla-profile.

If this parameter is configured within a sap-egress policy that is applied to any object except of the egress of an sla-profile, the configured parent schedulers will not be found and so the queues will not be parented and will be orphaned.

policer

Syntax **policer** *policer-id* [{[**port-redirect-group-queue**] [**queue** *queue-id*]} | {**group** *queue-group-name* [**instance** *instance-id*] [**queue** *group-queue-id]*}]
no policer

Context config>qos>sap-egress>fc

Description Within a sap-egress QoS policy forwarding class context, the policer command is used to map packets that match the forwarding class to the specified policer-id. The specified policer-id must already exist within the sap-egress QoS policy. The forwarding class of the packet is first discovered at ingress based on the ingress classification rules. When the packet arrives at egress, the sap-egress QoS policy may match a forwarding class reclassification rule which overrides the ingress derived forwarding class. The forwarding class context within the sap-egress QoS policy is then used to map the packet to an egress queue (using the queue queue-id, or port-redirect-group queue queue-id, or group queue-group-name instance instance-id queue queue-id commands) or an egress policer (policer policer-id). The queue and policer commands within the forwarding class context are mutually exclusive. By default, the forwarding class is mapped to the SAP egress default queue (queue 1). If the policer policer- id command is executed, any previous policer mapping or queue mapping for the forwarding class is overridden if the policer mapping is successful.

A policer defined within the sap-egress policy is not actually created on an egress SAP or a subscriber using an sla-profile where the policy is applied until at least one forwarding class is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or egress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class mapping will fail.

Packets that are mapped to an egress policer that are not discarded by the policer must be placed into a default queue on the packets destination port. The system uses egress port queue groups for this purpose. An egress queue group named policer-output-queues is automatically created on each port that support egress policers. By default, the system uses the forwarding class mappings within this queue group to decide which queue within the group will receive each packet output from the policer. This default policer output queuing behavior may be overridden for non-subscriber packets by redirection to a queue group. The name and instance of the queue group to redirect to is either spcecifed in the QoS policy itself, or the fact that a forwarding class must be redirected is simply identified in the QoS policy and the specific queue group instance is only identified at the time the QoS poicy is applied:

- If the **policer** *policer-id* command is successfully executed, the default egress queuing is performed for the forwarding class using the policer-output-queues queue group and the *queue-id* within the group based on the forwarding class map from the group template
- If the **policer** *policer-id* **queue** *queue-id* command is successfully executed, the specified SAP *queue-id* within egress QoS policy is used instead of the default policer output queues.
- If the **policer** *policer-id* **port-redirect-group-queue** keyword is successfully executed, the system will map the forwarding class to the queue within the egress queue group instance specified at the

time the QoS policy is applied to the SAP, using the forwarding class map from the queue group template.

- If the **policer** *policer-id* **port-redirect-group queue** *queue-id* command is successfully executed, the system will map the forwarding class to the configured *queue-id* within the egress queue group instance that is specified at the time the QoS policy is applied to the SAP (ignoring using the forwarding class map from the queue group template).
- If the **policer** *policer-id* **group** *queue-group-name* **instance** *instance-id* command is successfully executed, the system will map the forwarding class to the queue within the specified egress queue group instance using the forwarding class map from the group template.
- If the **policer** *policer-id* **group** *queue-group-name* **instance** *instance-id* **queue** *queue-id* command is successfully executed, the system will map the forwarding class to the specified *queue-id* within the specified egress queue group instance (ignoring the forwarding class map in the group template).

If the specified **group** *queue-group-name* is not defined as an egress queue-group-template, the **policer** command will fail. Further, if the specified group does not exist on the port for the SAPs or subscribers associated with the **sap-egress** QoS policy, the policer command will fail. While a group queue-group-name is specified in a **sap-egress** QoS policy, the groups corresponding egress template cannot be deleted. While a port egress queue group is associated with a policer instance, the port queue group cannot be deleted.

If the specified queue group-queue-id is not defined in the egress queue-group-template queue-group- name, the policer command will fail. While a *queue-id* within an egress queue group template is referenced by a **sap-egress** QoS policy forwarding class policer command, the queue cannot be deleted from the queue group template.

If an egress policed packet is discarded by the egress port queue group queue, the source policer discard stats are incremented. This means that the discard counters for the policer represent both the policer discard events and the destination queue drop tail events associated with the policer.

The **no** form of this command is used to restore the mapping of the forwarding class to the default queue. If all forwarding classes have been removed from the default queue, the queue will not exist on the SAPs or subscribers associated with the QoS policy and the no policer command will cause the system to attempt to create the default queue on each object. If the system cannot create the default queue in each instance, the no policer command will fail and the forwarding class will continue its mapping to the existing *policer-id*. If the no policer command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscribers will be lost.

Default none

Parameters *policer-id* — When the forwarding class **policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.

Values 1—8

Default none

port-redirect-group-queue — Used to override the forwarding class default egress queue destination to an egress port queue group. The specific egress queue group instance to use is specified at the time that the QoS policy is applied to the SAP. Therefore, this parameter os only valid if SAP based redirection is required.

queue *queue-id* **—** This parameter overrides the forwarding class default egress queue destination to a specified *queue-id*. If port-redirect-group is not configred, then this will be a local SAP queue of that *queue-id*. A queue of ID *queue-id* must exist within the egress QoS policy. If **port-redirect-group-queue** is configured then the the **queue** *queue-id* in the egress port queue group instance is used.

Values 1—8

Default Derived from forwarding class assignment in queue-group definition.

group *queue-group-name* **—** The **group** *queue-group-name* is optional and is used to override the forwarding class's default egress queue destination. If the queue group-queue-id parameter is not specified, the forwarding class map within the specified group's template is used to derive which queue within the group will receive the forwarding class's packets. An egress queue group template must exist for the specified queue-group-name or the policer command will fail. The specified queue-group-name must also exist as an egress queue group on the ports where SAPs and subscribers associated with the sap-egress policy is applied or the policer command will fail.

Values Any qualifying egress queue group name

Default **policer-output-queues**

queue *group-queue-id* — The **queue** *group-queue-id* is optional when the group queue-group-name parameter is specified and is used to override the forwarding class mapping within the group's egress queue group template. The specified group-queue-id must exist within the group's egress queue group template or the policer command will fail.

Values 1—8

Default Derived from forwarding class assignment in queue-group definition

instance *instance-id* **—** This parameter is used to specify the specific instance of a queue group with template queue-group-name to which this queue should be redirected. This parameter is only valid for queue groups on egress ports where policy based redirection is required.

Values 1 — 40960

Default 1

description

Syntax **description** *description string*
no description

Context config>qos>sap-egress>policer

Description The **description** command is used to define an informational ASCII string associated with the policer control policy. The string value can be defined or changed at any time once the policy exists.

The **no** form of this command is used to remove an explicit description string from the policer.

Default **no description**

Parameters *description string* **—** The *description-string* parameter defines the ASCII description string for the policer control policy. The description string can be up to 80 characters. If the string contains spaces, it must be placed within beginning and ending double quotation marks. Beginning and ending quotation marks are

not considered part of the description string. Only printable ASCII characters are allowed in the string. The sting does not need to be unique and may be repeated in the descriptions for other policer control policies or other objects. If the command is executed without the description-sting present, any existing description string will be unaffected.

Default None

adaptation-rule

Syntax **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
no adaptation-rule

Context config>qos>sap-egress>policer

Description This command is used to define how the policer's configuration parameters are translated into the underlying hardware capabilities used to implement each policer instance. For instance, the configured rates for the policer need to be mapped to the timers and decrement granularity used by the hardware's leaky bucket functions that actually perform the traffic metering. If a rate is defined that cannot be exactly matched by the hardware, the adaptation-rule setting provides guidance for which hardware rate should be used.

The **max** keyword tells the system that the defined rate is the maximum rate that should be given to the policer. If the hardware cannot exactly match the given rate, the next lowest hardware supported rate is used.

The **min** keyword tells the system that the defined rate is the minimum rate that should be given to the policer. If the hardware cannot exactly match the given rate, the next highest hardware supported rate is used.

The **closest** keyword tells the system that the defined rate is the target rate for the policer. If the hardware cannot exactly match the given rate, the system will use the closest hardware supported rate compared to the target rate.

The hardware also needs to adapt the given mbs and cbs values into the PIR bucket violate threshold (discard) and the CIR bucket exceed threshold (out-of-profile). The hardware may not have an exact threshold match which it can use. In R8.0, the system treats the mbs and cbs values as minimum threshold values.

The **no** form of this command is used to return the policer's metering and profiling hardware adaptation rules to closest.

Parameters **pir** {**max** | **min** | **closest**} — When the optional **pir** parameter is specified, the **max**, **min** or **closest** keyword qualifier must follow.

max — The **max** keyword is used to inform the system that the metering rate defined for the policer is the maximum allowed rate. The system will choose a hardware supported rate that is closest but not exceeding the specified rate.

min — The **min** keyword is used to inform the system that the metering rate defined for the policer is the minimum allowed rate. The system will choose a hardware supported rate that is closest but not lower than the specified rate.

closest — The **closest** keyword is used to inform the system that the metering rate defined for the policer is the target rate. The system will choose a hardware supported rate that is closest to the specified rate.

Default closest

cir {**max** | **min** | **closest**} — When the optional **cir** parameter is specified, the **max**, **min** or **closest** keyword qualifier must follow.

max — The **max** keyword is used to inform the system that the profiling rate defined for the policer is the maximum allowed rate. The system will choose a hardware supported rate that is closest but not exceeding the specified rate.

min — The min keyword is used to inform the system that the profiling rate defined for the policer is the minimum allowed rate. The system will choose a hardware supported rate that is closest but not lower than the specified rate.

closest — The closest keyword is used to inform the system that the profiling rate defined for the policer is the target rate. The system will choose a hardware supported rate that is closest to the specified rate.

Default closest

cbs

Syntax **cbs** {*size* [**bytes** | **kilobytes**] | **default**}
no cbs

Context config>qos>sap-egress>policer

Description This command is used to configure the policer's CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the policer. If the policer's forwarding rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.

The policer's **cbs** size defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The no form of this command returns the policer to its default CBS size.

Default 64 kilobytes when CIR = max, otherwise, 10ms volume of traffic for a configured non zero/non max CIR.

Parameters *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **cbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

byte — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobyte — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Values 0 – 16777216 or **default**

Default **kilobyte**

high-prio-only

Syntax **high-prio-only** *percent-of-mbs*
no high-prio-only

Context config>qos>sap-egress>policer

Description This command is used to configure the percentage of the policer's PIR leaky bucket's MBS (maximum burst size) that is reserved for high priority traffic. While the **mbs** value defines the policer's high priority violate threshold, the percentage value defined is applied to the **mbs** value to derive the bucket's low priority violate threshold. See the **mbs** command details for information on which types of traffic is associated with each violate threshold.

Default **high-prio-only 10**

Parameters *percent-of-mbs* — The *percent-of-mbs* parameter is required when specifying **high-prio-only** and is expressed as a percentage with granularity of 1,000th of a percent.

Values 0—100

Default 10

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
no mbs

Context config>qos>sap-egress>policer

Description This command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low priority violate threshold. For egress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold.

The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by **high-prio-only** is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.

The no form of this command returns the policer to its default MBS size.

Default None

Parameters *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

byte — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobyte — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Values 0 – 16777216 or **default**

Default **kilobyte**

packet-byte-offset

Syntax **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}
no packet-byte-offset

Context config>qos>sap-egress>policer

Description This command is used to modify the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** version of this command is used to remove per packet size modifications from the policer.

Parameters **add** *bytes* — The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 0—31

Default None

subtract *bytes* — The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When b is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 0—31

Default None

parent

Syntax **parent** {**root** | *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]
no parent

Context config>qos>sap-egress>policer

Description This command is used to create a child to parent mapping between each instance of the policer and either the **root** arbiter or a specific tiered **arbiter** on the object where the policy is applied. Defining a parent association for the policer causes the policer to compete for the parent policer's available bandwidth with other child policers mapped to the policer control hierarchy.

Policer control hierarchies may be created on SAPs or on a subscriber context. To create a policer control hierarchy on an ingress or egress SAP context, a **policer-control-policy** must be applied to the SAP. Once applied, the system will create a parent policer that is bandwidth limited by the policy's **max-rate** value under the root arbiter. The root arbiter in the policy also provides the information used to determine the various priority level discard-unfair and discard-all thresholds. Besides the root arbiter, the policy may also contain user defined tiered arbiters that provide arbitrary bandwidth control for subsets of child policers that are either directly or indirectly parented by the arbiter.

When the QoS policy containing the policer with a **parent** mapping to an arbiter name exists on the SAP, the system will scan the available arbiters on the SAP. If an arbiter exists with the appropriate name, the policer to arbiter association is created. If the specified arbiter does not exist either because a **policer-control-policy** is not currently applied to the SAP or the arbiter name does not exist within the applied policy, the policer is placed in an orphan state. Orphan policers operate as if they are not parented and are not subject to any bandwidth constraints other than their own PIR. When a policer enters the orphan state, it is flagged as operationally degraded due to the fact that it is not operating as intended and a trap is generated. Whenever a **policer-control-policy** is added to the SAP or the existing policy is modified, the SAP's policer's parenting configurations must be reevaluated. If an orphan policer becomes parented, the degraded flag should be cleared and a resulting trap should be generated.

For subscribers, the policer control hierarchy is created through the **policer-control-policy** applied to the **sub-profile** used by the subscriber. A unique policer control hierarchy is created for each subscriber associated with the **sub-profile**. The QoS policy containing the policer with the parenting command comes into play through the subscriber **sla-profile** which references the QoS policy. The combining of the **sub-profile** and the **sla-profile** at the subscriber level provides the system with the proper information to create the policer control hierarchy instance for the subscriber.

Executing the **parent** command will fail if:

- The policer's stat-mode in the QoS policy is set to no-stats
- A stat-mode no-stats override exists on an instance of the policer on a SAP or subscribers context

A policer with a **parent** command applied cannot be configured with **stat-mode no-stats** in either the QoS policy or as an override on an instance of the policer

Once a policer is successfully parented to an arbiter, the **parent** commands **level** and **weight** parameters are used to determine at what priority level and at which weight in the priority level that the child policer competes with other children (policers or other arbiters) for bandwidth.

The **no** form of this command is used to remove the parent association from all instances of the policer.

Parameters {**root** | *arbiter-name*} — When the **parent** command is executed, either the keyword **root** or an *arbiter-name* must be specified.

root — The **root** keyword specifies that the policer is intended to become a child to the **root** arbiter where an instance of the policer is created. If the **root** arbiter does not exist, the policer will be placed in the orphan state.

arbiter-name — The *arbiter-name* parameter specifies that the policer is intended to become a child to one of the tiered arbiters with the given arbiter-name where an instance of the policer is created. If the specified arbiter-name does not exist, the policer will be placed in the orphan state.

Default None

weight *weight-within-level* — The **weight** *weight-within-level* keyword and parameter are optional when executing the **parent** command. When **weight** is not specified, a default level of 1 is used in the parent arbiters priority level. When **weight** is specified, the *weight-within-level* parameter must be specified as an integer value from 1 through 100.

Default 1

rate

Syntax **rate** {**max** | **kilobits-per-second**} [**cir** {**max** | **kilobits-per-second**}]
no rate

Context config>qos>sap-egress>policer

Description This command is used to configure the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on the each packets size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches it's exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs** and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow based on the conforming or exceeding state from the CIR bucket.

When a packet is red neither the PIR or CIR bucket depths are incremented by the packets size. When the packet is yellow the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 Kbps (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command is used to restore the default metering and profiling rate to a policer.

Parameters {**max** | *kilobits-per-second*} — Specifying the keyword **max** or an explicit *kilobits-per-second* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to max.

Values **max** or 1—2000000000

Values **max** or 0—2000000000

cir {**max** | *kilobits-per-second*} — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *kilobits-per-second* parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 Kbps. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second.

Values **max** or 0—20,000,000

stat-mode

Syntax **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-profile-cir** | **offered-total-cir** | **offered-profile-capped-cir** | **offered-limited-capped-cir**}
no stat mode

Context config>qos>sap-egress>policer
config>qos>queue-group-templates>egress>queue-group

Description The **sap-egress** QoS policy's policer **stat-mode** command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. An egress policer has multiple types of offered packets (soft in-profile and out-of-profile from ingress and hard in-profile and out-of-profile due to egress profile overides) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the potential large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly re-profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported which prevents any packet accounting, the use of the policer's **parent** command requires at the policer's **stat-mode** to be set at least to the **minimal** setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. Once a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding

plane's policer counter resources. You can view the the total/allocated/free stats by using the **tools dump system-resources** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The default **stat-mode** when a policer is created within the policy is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's **stat-mode** setting to **minimal**. The command will fail if insufficient policer counter resources exist to implement **minimal** where the QoS policer is currently applied and has a forwarding class mapping.

Parameters **no-stats —** Counter resource allocation:0

The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using **no-stats** cannot be a child to a parent policer and the policer's **parent** command will fail.

When **collect-stats** is enabled, the lack of counters causes the system to generate the following statistics:

a. offered-in = 0

b. offered-out = 0

c. discard-in = 0

d. discard-out = 0

e. forward-in = 0

f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

minimal — Counter resource allocation:1

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (profile or priority) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters are used in the following manner:

1. offered = soft-in-profile-out-of-profile, profile in/out

2. discarded = Same as 1

3. forwarded = Derived from 1 - 2

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out = 0

c. discard-in = 2

d. discard-out = 0

e. forward-in = 3

f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

offered-profile-no-cir — Counter resource allocation:2

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when profile based offered, discard and forwarding stats are required from the egress policer, but a CIR is not being used to recolor the soft in-profile and out-of-profile packets. This mode does not prevent the policer from being configured with a CIR rate.

The counters are used in the following manne:

1. offered-in = soft-in-profile, profile in

2. offered-out = soft-out-of-profile, profile out

3. dropped-in = Same as 1

4. dropped-out = Same as 2

5. forwarded-in = Derived from 1 - 3

6. forwarded-out = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out = 2

c. discard-in = 3

d. discard-out = 4

e. forward-in = 5

f. forward-out = 6

offered-profile-cir — Counter resource allocation: 3

The **offered-profile-cir** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-profile-cir** mode is most useful when profile based offered, discard and forwarding stats are required from the egress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

The counters are used in the following manner:

1. offered-in-that-stayed-green-or-turned-red = profile in

2. offered-soft-that-turned-green = soft-in-profile-out-of-profile

3. offered-soft-or-out-that-turned-yellow-or-red = soft-in-profile-out-of-profile, profile out

4. dropped-in-that-stayed-green-or-turned-red = Same as 1

5. dropped-soft-that-turned-green = Same as 2

6. dropped-soft-or-out-that-turned-yellow-or-red = Same as 3

7. forwarded-in-that-stayed-green = Derived from 1 - 4

8. forwarded-soft-that-turned-green = Derived from 2 - 5

9. forwarded-soft-or-out-that-turned-yellow = Derived from 3 - 6

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out = 2 + 3

c. discard-in = 4

d. discard-out = 5 + 6

e. forward-in = 7 + 8

f. forward-out = 9

offered-total-cir — Counter resource allocation:2

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high and low priority classifications are not being used on the un-trusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the un-trusted packets.

The counters are used in the following manner:

1. offered-that-turned-green =soft-in-profile-out-of-profile, profile in/out

2. offered- that-turned-yellow-or-red =soft-in-profile-out-of-profile, profile in/out

3. dropped-offered-that-turned-green = Same as 1

4. dropped-offered-that-turned-yellow-or-red = Same as 2

5. forwarded-offered-that-turned-green = Derived from 1 - 3

6. forwarded-offered-that-turned-yellow = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1 + 2

b. offered-out = 0

c. discard-in = 3

d. discard-out = 4

e. forward-in = 5

f. forward-out = 6

Counter 0 indicates that the accounting statistic returns a value of zero.

offered-profile-capped-cir — Counter resource allocation:2

When **offered-profile-capped-cir** is defined, the system creates five offered-output counters in the forwarding plane and five discard counters in the traffic manager.

The **offered-profile-capped-cir** mode is similar to the **offered-profile-cir** mode except that it includes support for **profile in** and **soft-in-profile** that may be output as 'out-of-profile' due to enabling profile-capped mode on the ingress policer.

The impact of using o**ffered-profile-capped-cir** stat-mode while **profile-capped** mode is disabled are that one of the counting resources in the forwarding plane and traffic manager will not be used and soft-in-profile will be treated as 'offered-in' instead of 'offered-undefined'.

The counters are used in the following manner:

1. 'offered-in-that-stayed-green' = profile in, soft-in-profile
2. 'offered-in-that-turned-yellow-or-red' = profile in, soft-in-profile
3. 'offered-soft-out-that-turned-green' = soft-out-of-profile
4. 'offered-soft-out- that-turned-yellow-or-red' = soft-out-of-profile
5. 'offered-out-that-turned-yellow-or-red' = profile out
6. 'dropped-in-that-stayed-green' = Same as 1
7. 'dropped-in-that-turned-yellow-or-red' = Same as 2
8. 'dropped-soft-out-that-turned-green' = Same as 3
9. 'dropped-soft-out-that-turned-yellow-or-red' = Same as 4
10. 'dropped-out-that-turned-yellow-or-red' = Same as 5
11. 'forwarded-in-that-stayed-green' = Derived from 1 - 6
12. 'forwarded-in-that-turned-yellow' = Derived from 2 - 7
13. 'forwarded-soft-out-that-turned-green' = Derived from 3 - 8
14. 'forwarded-soft-out-that-turned-yellow' = Derived from 4 - 9
15. 'forwarded-out-that-turned-yellow' = Derived from 5 - 10

When c**ollect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. 'offered-undefined' = 3 + 4

b. 'offered-in' = 1 + 2

c. 'offered-out' = 5

d. 'discard-in' = 6 + 8

e. 'discard-out' = 7 + 9 + 10

f. 'forward-in' = 11 + 13

g. ‘forward-out’ = 12 + 14 + 15

offered-limited-capped-cir — Counter resource allocation:2

offered-limited-capped-cir is defined, the system creates four forwarding plane offered-output counters in the network processor and three discard counters in the traffic manager.

The **offered-limited-capped-cir** mode is similar to the **offered-profile-capped-cir** mode except that it combines **profile out** and **soft-out-of-profile** and eliminates the ‘offered-undefined’ statistic.

The impact of using **offered-limited-capped-cir** stat-mode while profile-capped mode is disabled are that one of the counting resources in the forwarding plane and traffic manager will not be used and **soft-in-profile** will be treated as ‘offered-in’ instead of ‘offered-undefined’.

The counters are used in the following manner:

1. ‘offered-in-that-stayed-green’ = profile in, soft-in-profile
2. ‘offered-in-that-turned-yellow-or-red’ = profile in, soft-in-profile
3. ‘offered-out-that-turned-green’ = soft-out-of-profile
4. ‘offered-out- that-turned-yellow-or-red’ = profile out, soft-out-of-profile
5. ‘dropped-in-that-stayed-green’ = Same as 1
6. ‘dropped-in-that-turned-yellow-or-red’ = Same as 2
7. ‘dropped-out-that-turned-green’ = Same as 3
8. ‘dropped-out-that-turned-yellow-or-red’ = Same as 4
9. ‘forwarded-in-that-stayed-green’ = Derived from 1 - 5
10. ‘forwarded-in-that-turned-yellow’ = Derived from 2 - 6
11. ‘forwarded-out-that-turned-green’ = Derived from 3 - 7
12. ‘forwarded-out-that-turned-yellow’ = Derived from 4 – 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. ‘offered-in’ = 1 + 2
b. ‘offered-out’ = 3 + 4
c. ‘discard-in’ = 5 + 7
d. ‘discard-out’ = 6 + 8
e. ‘forward-in’ = 9 + 11
f. ‘forward-out’ = 10 + 12

dscp

Syntax **dscp** {*dscp-name* | **in-profile** *dscp-name* **out-profile** *dscp-name*}
no dscp

Context config>qos>sap-egress>fc

Description This command configures a DiffServ Code Point (DSCP) code point to be used for remarking packets from the specified FC. If the optional in/out-profile is specified, the command will remark different DSCP code points depending on whether the packet was classified to be in or out-of-profile ingress to the node.

Default not enabled

Parameters *dscp-name* — Specifies a dscp name that has been previously mapped to a value using the **dscp-name** command. The DiffServ code point can only be specified by its name.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

prec

Syntax **prec** *ip-prec-value* [**fc** *fc-name*] [**hsmda-counter-override** *counter-id*] [**profile** {**in** | **out**}]
no prec *ip-prec-value*

Context config>qos>sap-egress>fc

Description This command defines a value to be used for remarking packets for the specified FC. If the optional in/out-profile is specified, the command will remark different PREC values depending on whether the packet was classified to be in or out-of-profile ingress to the node.

The hsmda-counter-override keyword is optional. When specified and the egress SAP is created on an HSMDA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMDA uses each queues default queue counters for packets mapped to the queue. The hsmda-counter-override keyword is used to map the packet to an explicit exception counter. One of eight counters may be used. When the packet is mapped to an exception counter, the packet will not increment the queues discard or forwarding counters, instead the exception discard and forwarding counters will be used. The dscp based counter override decision may be overwritten by an ip-criteria reclassification rule match if the higher priority classification rule has an hsmda-counter-override action defined.

not enabled

Parameters *ip-prec-value* — The *ip-prec-value* is a required parameter that specifies the unique IP header ToS byte precedence bits value that will match the IP precedence rule. If the command is executed more than once with the same *ip-prec-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight IP precedence rules are allowed on a single policy.

The precedence is evaluated from the lowest to highest value.

Values 0 — 7

hsmda-counter-override *counter-id* — The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the IP precedence value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The hsmda-counter-override action may be overwritten by an ip-criteria

reclassification rule match. The specified counter-id must be specified as an integer between 1 and 8. To remove the ESMDA exception counter reclassification action for the specified dscp-value, the dscp command must be re-executed without the hsmda-counter-override reclassification action defined.

Values 1 — 8

scope

Syntax **scope** {**exclusive** | **template**}
no scope

Context config>qos>sap-egress *policy-id*

Description Enter the scope of this policy. The scope of the policy cannot be changed if the policy is applied to one or more services.

The no form of this command sets the scope of the policy to the default of template.

Default template

Parameters **exclusive —** When the scope of a policy is defined as exclusive, the policy can only be applied to a single SAP. Attempting to assign the policy to a second SAP will result in an error message. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.

The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1.

template — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs on the router.

sap-egress

Syntax [**no**] **sap-egress** *policy-id* | *policy-name*

Context config>qos

Description This command is used to create or edit a Service Egress QoS policy. The egress policy defines the Service Level Agreement (SLA) for service packets as they egress on the SAP.

Policies in effect are templates that can be applied to multiple services as long as the scope of the policy is template. The queues defined in the policy are not instantiated until a policy is applied to a service.

A sap-egress policy differs from sap-ingress policies in the complexity of the QoS parameters that can be

defined. At ingress, policies determine queue mappings based on ingress DSCP, Dot1P and IP or MAC match criteria. Multiple queues can be created per forwarding class and each queue can have different CIR or PIR parameters.

At egress, the policies are much simpler, as the forwarding class and in or out of profile determination happened way back at the original service ingress SAP. Egress SAP QoS policies allow the definition of queues and the mapping of forwarding classes to those queues. Each queue needs to have a relative CIR for determining its allocation of QoS resources during periods of congestion. A PIR can also be defined that forces a hard limit on the packets transmitted through the queue. When the forwarding class is mapped to the queue, a Dot1p value can optionally be specified. If specified and the SAP has a Dot1q encapsulation type, the Dot1p value will be used for all packets that egress on that forwarding class. If the Dot1p value is not specified, a Dot1p value of zero will be used. If the SAP is null encapsulated, or on a SONET/SDH interface, the Dot1p value has no meaning.

A **default-action** parameter is required to specify the default queue used by all forwarding classes not specifically mapped within the queue parameters. A sap-egress policy will be considered incomplete, if it does not include definition of at least one queue and does not specify the default action. Incomplete sap-egress policies cannot be applied to services.

The sap-egress policy with policy-id 1 is the default sap-egress QoS policy and is applied to service egress SAPs when an explicit policy is not specified or removed. The system sap-egress policy can be modified but not deleted. Using the **no sap-egress** command on **policy-id 1** causes it to revert to its factory default parameters.

The factory default settings for sap-egress policy-id 1 define a single queue with PIR set to the maximum value and a CIR set to 25. The single queue is the default queue and all forwarding classes will map to it. Packets being tagged according to the SAP encapsulation defined will have the Dot1p bits set to zero.

Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all egress SAPs where this policy is applied. For this reason, when many changes are required on a policy, it is highly recommended that the policy be copied to a work area policy-id. That work-in-progress policy can be modified until complete and then written over the original policy-id. Use the **config qos copy** command to maintain policies in this manner.

The no form of this command to deletes the sap-egress policy. A policy cannot be deleted until it is removed from all service SAPs where it is applied. When a sap-egress policy is removed from a SAP, the SAP will revert to the default sap-egress policy-id 1.

The system default sap-egress policy is a special case. The **no** command restores the factory defaults to policy-id 1.

Parameters *policy-id* — The policy-id uniquely identifies the policy on the router.

Default none

Values 1 — 65535

policy-name — The *policy-name* uniquely identifies the policy.

Values Valid names consist of any string up to 64 characters long. Policies must first be created with a policy-id, after which a policy-name can be assigned and used as an alias to reference the policy during configuration changes. Policy names may not begin with a number (0-9) or the underscore "_" character (e.g. _myPolicy). "Default" or "default" can not be used as policy names. Saved configurations and display output from the "info"

and most "show" commands will show the policy-id (not the policy-name) where the policies are referenced.

de-mark

Syntax [**no**] **de-mark** [**force** *de-value*]

Context config>qos>sap-egress>fc

Description This command is used to explicitly define the marking of the DE bit for **fc** *fc-name* according to the in and out of profile status of the packet (fc-name may be used to identify the dot1p-value).

If no de-value is present, the default values are used for the marking of the DE bit: for example, 0 for in-profile packets, 1 for out-of-profile ones– see IEEE 802.1ad-2005 standard.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS/PBB-Epipe, the command dictates the marking of the DE bit for both the BVID and ITAG.

If this command is not used, the DE bit should be preserved if an ingress TAG exist or set to zero otherwise.

If the de-value is specifically mentioned in the command line it means this value is to be used for all the packets of this forwarding class regardless of their in/out of profile status.

The commands **de-mark-inner** and **de-mark-outer** take precedence over the **de-mark** command if both are specified in the same policy.

Values 0 or 1

dot1p

Syntax [**no**] **dot1p** {*dot1p-value* | **in-profile** *dot1p-value* **out-profile** *dot1p-value*}

Context config>qos>sap-egress>fc *fc-name*

Description This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the dot1p command has no effect.

The optional **in-profile** | *dot1p-value* **out-profile** *dot1p-value* structure added to the existing dot1p command will add the capability to mark on an egress SAP the in and out of profile status via a certain dot1p combination, similarly with the DE options.

The command with the additional structure may be used on the SAP when the internal in and out of profile status needs to be communicated to an access network/customer device that does not support the DE bit. Once the in-profile keyword is added, then the rest of the newly added structure must be specified.

When these commands are used the DE Bit or the equivalent field is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DE bit is set to 0.

When the previous command (dot1p dot1p-value) is used without the new structure, it means that the dot1p-value is used for the entire forwarding class, same as before. The two versions of the command are mutually exclusive.

Independently the in or out profile status may be indicated via the setting of the DE bit setting if the de-mark command is used.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS/PBB-Epipe, the command dictates the marking of the dot1p bits for both the BVID and ITAG.

The commands **dot1p-inner** and **dot1p-outer** take precedence over the dot1p command if both are specified in the same policy.

The **no** form of the command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.

Default 0

in-profile *dot1p-value* **—** Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 — 7

out-profile *dot1p-value* **—** specifies the 802.1p value to set for out-profile frames in this forwarding class.

Values 0 — 7

de-mark-inner

Syntax **[no] de-mark-inner [force de-value]**

Context config>qos>sap-egress>fc

Description This command is used to explicitly define the marking of the DE bit in the inner VLAN tag for fc fc-name on a qinq SAP according to the in and out of profile status of the packet.

If no de-value is present, the default values are used for the marking of the DE bit: for example, 0 for in-profile packets, 1 for out-of-profile ones– see IEEE 802.1ad-2005 standard.

If the de-value is included in the command line then this value is used for all the inner tags of packets of this forwarding class regardless of their in/out of profile status.

This command takes precedence over the **de-mark** command if both are specified in the same policy and over the default action.

The configuration of **qinq-mark-top-only** under the SAP egress takes precedence over the use of the **de-mark-inner** in the policy, i.e. the inner VLAN tag is not remarked when **qinq-mark-top-only** is configured (the marking used for the inner VLAN tag is based on the current default which is governed by the marking of the packet received at the ingress to the system).

If **no de-mark** commands are used, the DE bit is preserved if an ingress inner tag exists, or set to zero otherwise.

This command is only supported on FP2 and higher based hardware, and is otherwise ignored.

Values 0 or 1

de-mark-outer

Syntax **[no] de-mark-outer [force** *de-value***]**

Context config>qos>sap-egress>fc

Description This command is used to explicitly define the marking of the DE bit in the outer or single VLAN tag on a qinq or dot1q SAP, respectively, according to the in and out of profile status of the packet.

If no de-value is present, the default values are used for the marking of the DE bit: for example, 0 for in-profile packets, 1 for out-of-profile ones– see IEEE 802.1ad-2005 standard.

If the de-value is included in the command line then this value is used for all the outer or single tags of packets of this forwarding class regardless of their in/out of profile status.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS/PBB-Epipe, the command dictates the marking of the DE bit for both the BVID and ITAG.

This command takes precedence over the **de-mark** command if both are specified in the same policy and over the default action.

If **no de-mark** commands are used, the DE bit is preserved if an ingress outer or single tag exists, or set to zero otherwise.

This command is supported on FP2 and higher based hardware, and is otherwise ignored.

Values 0 or 1

dot1p-inner

Syntax **[no] dot1p-inner {***dot1p-value* **| in-profile** *dot1p-value* **out-profile** *dot1p-value***}**

Context config>qos>sap-egress>fc

Description This command explicitly defines the egress inner VLAN tag IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets of *fc-name* that have either an inner IEEE 802.1Q or IEEE 802.1P encapsulation on a qinq SAP will use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P qinq encapsulated, this command has no effect.

The optional **in-profile** | *dot1p-value* **out-profile** *dot1p-value* parameters on the **dot1p-inner** command adds the capability to mark the in and out of profile status on an egress qinq SAP. The command with the additional parameters may be used on the SAP when the internal in and out of profile status needs to be communicated to an access network/customer device that does not support the DE bit. Once the in-profile keyword is added, then the rest of the structure must be specified.

When these commands are used, the DE Bit or the equivalent field is left unchanged by the egress processing if an inner tag exists. If a new inner tag is added, the related DE bit is set to 0. The in or out profile status may be indicated via the setting of the DE bit setting if the **de-mark(-inner)** command is used.

The two versions of the command (with and without parameters) are mutually exclusive.

This command takes precedence over the **dot1p** command if both are specified in the same policy and over the default action where the marking is taken from packet received at ingress.

The configuration of **qinq-mark-top-only** under the SAP egress takes precedence over the use of the

dot1p-inner in the policy, that is, the inner VLAN tag is not remarked when **qinq-mark-top-only** is configured (the marking used for the inner VLAN tag is based on the current default which is governed by the marking of the packet received at the ingress to the system).

The **no** form of the command sets the inner IEEE 802.1P or IEEE 802.1Q priority bits to 0.

This command is supported on FP2 and higher based hardware, and is otherwise ignored.

Default 0

Parameters **dot1p-inner** *dot1p-value* **—** Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 — 7

in-profile *dot1p-value* **—** Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 — 7

out-profile *dot1p-value* **—** specifies the 802.1p value to set for out-profile frames in this forwarding class.

Values 0 — 7

dot1p-outer

Syntax **[no] dot1p-outer {***dot1p-value* **| in-profile** *dot1p-value* **out-profile** *dot1p-value***}**

Context config>qos>sap-egress>fc

Description This command explicitly defines the egress outer or single VLAN tag IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets of fc-name that have either an outer or single IEEE 802.1Q or IEEE 802.1P encapsulation on a qinq or a dot1p SAP, respectively, will use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, this command has no effect.

The optional **in-profile** | *dot1p-value* **out-profile** *dot1p-value* parameters on the dot1p-outer command adds the capability to mark the in and out of profile status on an egress qinq or dot1p SAP. The command with the additional parameters may be used on the SAP when the internal in and out of profile status needs to be communicated to an access network/customer device that does not support the DE bit. Once the in-profile keyword is added, then the rest of the structure must be specified.

When these commands are used, the DE Bit or the equivalent field is left unchanged by the egress processing if a (single or outer) tag exists. If a new tag is added, the related DE bit is set to 0. The in or out profile status may be indicated via the setting of the DE bit setting if the **de-mark(-outer)** command is used.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS/PBB-Epipe, the command dictates the marking of the dot1p bits for both the BVID and ITAG.

The two versions of the command (with and without parameters) are mutually exclusive.

This command takes precedence over the **dot1p** command if both are specified in the same policy and over the default action where the marking is taken from packet received at ingress.

The **no** form of the command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.

This command is supported on FP2 and higher based hardware, and is otherwise ignored.

Default 0

Parameters **dot1p-outer** *dot1p-value* — Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 — 7

in-profile *dot1p-value* — Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 — 7

out-profile *dot1p-value* — specifies the 802.1p value to set for out-profile frames in this forwarding class.

Values 0 — 7

Service Queue QoS Policy Commands

adaptation-rule

Syntax **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
no adaptation-rule

Context config>qos>sap-ingress>queue
config>qos>sap-egress>queue

This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default adaptation-rule pir closest cir closest

Parameters *adaptation-rule* — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.

Values **pir** — Defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* rate command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — Defines the constraints enforced when adapting the CIR rate defined within the **queue** queue-id **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

max — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

adv-config-policy

Syntax [**no**] **adv-config-policy** *policy-name*

Context config>qos>sap-ingress>queue
config>qos>sap-egress>queue

Description This command specifies the advanced QoS policy. The advanced QoS policy contains only queue and policer child control parameters within a child-control node.

Once a policy is created, it may be applied to a queue or policer defined within a **sap-egress** or **sap-ingress** QoS policy. It may also be applied to a queue or policer defined within an ingress or **egress queue-group** template. When a policy is currently associated with a QoS policy or template, the policy may be modified but not deleted (even in the event that the QoS policy or template is not in use).

The **no** form of this command removes the specified advanced policy.

Default None

Parameters *policy-name* — The name of the advanced QoS policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

Context config>qos>sap-egress>queue

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-Load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet-based offered-load.
- Frame-encapsulation overhead — Using the avg-frame-overhead parameter, the frame-encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10,000 octets and the avg-frame-overhead equals 10%, the frame-encapsulation overhead would be 10,000 x 0.1 or 1,000 octets.

For egress Ethernet queues, the frame-encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame-encapsulation overhead would be 50 x 20 or 1,000 octets.

- Frame-based offered-load — The frame-based offered-load is calculated by adding the offered-load to the frame-encapsulation overhead. If the offered-load is 10,000 octets and the encapsulation overhead was 1,000 octets, the frame-based offered-load would equal 11,000 octets.
- Packet to frame factor — The packet -to-frame factor is calculated by dividing the frame-encapsulation overhead by the queue's offered-load (packet based). If the frame-encapsulation overhead is 1,000 octets and the offered-load is 10,000 octets then the packet to frame factor would be 1,000 / 10,000 or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- Frame-based CIR — The frame-based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame-based CIR would be 500 x 1.1 or 550 octets.
- Frame-based within-cir offered-load — The frame-based within-cir offered-load is the portion of the frame-based offered-load considered to be within the frame-based CIR. The frame-based within-cir offered-load is the lesser of the frame-based offered-load and the frame-based CIR. If the frame-based offered-load equaled 11000 octets and the frame-based CIR equaled 550 octets, the frame-based within-cir offered-load would be limited to 550 octets. If the frame-based offered-load equaled 450 octets and the frame-based CIR equaled 550 octets, the frame-based within-cir offered-load would equal 450 octets (or the entire frame-based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame-based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- Frame-based PIR — The frame-based PIR is calculated by multiplying the packet to frame-factor with the queue's-configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame-factor equals 0.1, the frame-based PIR would be 7,500 x 1.1 or 8,250 octets.
- Frame-based within-pir offered-load — The frame-based within-pir offered-load is the portion of the frame-based offered-load considered to be within the frame-based PIR. The frame-based within-pir offered-load is the lesser of the frame-based offered-load and the frame-based PIR. If the frame-based offered-load equaled 11,000 octets and the frame-based PIR equaled 8250 octets, the frame-based within-pir offered-load would be limited to 8,250 octets. If the frame-based offered-load equaled 7,000 octets and the frame-based PIR equaled 8,250 octets, the frame-based within-pir offered load would equal 7,000 octets.

Port Scheduler Operation Using Frame Transformed Rates — The port scheduler uses the frame based rates to figure the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and Subscriber SLA-Profile Average Frame Overhead Override — The average frame overhead parameter on a sap-egress may be overridden on an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers. An avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the

average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default 0

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0.00 — 100.00

burst-limit

Syntax **burst-limit** {**default** | *size* [**byte** | **kilobyte**]}
no burst-limit

Context config>qos>sap-ingress>queue
config>qos>sap-egress>queue

Description The **queue burst-limit** command is used to define an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The **burst-limit** command is supported under the sap-ingress and sap-egress QoS policy queues. The command is also supported under the ingress and egress queue-group-templates queues.

The **no** form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies or queue group templates. When specified within a queue-override queue context, any current burst limit override for the queue will be removed and the queue's burst limit will be controlled by its defining policy or template.

Parameters **default —** The default parameter is mutually exclusive to specifying an explicit size value. When burst-limit default is executed, the queue is returned to the system default value.

size — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in Kilobytes. If the value is intended to be interpreted in bytes, the byte qualifier must be added following size.

Values 1 to 14,000 (14,000 or 14,000,000 depending on bytes or kilobytes)

Default No default for size, use the default keyword to specify default burst limit

byte — The **bytes** qualifier is used to specify that the value given for size must be interpreted as the burst limit in bytes. The byte qualifier is optional and mutually exclusive with the kilobytes qualifier.

kilobyte — The **kilobyte** qualifier is used to specify that the value given for size must be interpreted as the burst limit in Kilobytes. The kilobyte qualifier is optional and mutually exclusive with the bytes qualifier. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.

burst-limit

Syntax **burst-limit** *size* [**bytes**|**kilobytes**]
no burst-limit

Context config>qos>sap-ingress>hsmda-queue>queue
config>qos>sap-egress>hsmda-queue>queue

Description The **queue burst-limit** command is used to define an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The **burst-limit** command is supported under the sap-ingress and sap-egress QoS policy queues. The command is also supported under the ingress and egress queue-group-templates queues.

The **no** form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies or queue group templates. When specified within a queue-override queue context, any current burst limit override for the queue will be removed and the queue's burst limit will be controlled by its defining policy or template.

Default no burst-limit

Parameters *size* — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in Kilobytes. If the value is intended to be interpreted in bytes, the byte qualifier must be added following size.

Values 1 to 1000000

Default No default for size, use the default keyword to specify default burst limit

byte — The **bytes** qualifier is used to specify that the value given for size must be interpreted as the burst limit in bytes. The byte qualifier is optional and mutually exclusive with the kilobytes qualifier.

kilobyte — The **kilobyte** qualifier is used to specify that the value given for size must be interpreted as the burst limit in Kilobytes. The kilobyte qualifier is optional and mutually exclusive with the bytes qualifier. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>qos>sap-egress>queue
config>qos>sap-ingress>queue

Description This command provides a mechanism to override the default reserved buffers for the queue. It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potentially large number of service queues and the economy of statistical multiplexing the individual queue's CBS settings into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS

buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly dropping packets.

If the CBS value is larger than the MBS value, the CBS is capped to the value of the MBS or the minimum CBS value. If the MBS and CBS values are configured to be equal (or nearly equal) this will result in the CBS being slightly higher than the value configured.

The **no** form of this command returns the CBS size to the default value.

Default default

Parameters *size-in-kbytes* — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values

0 — 131072 or default	
Minimum configurable non-zero value	6Kbytes
Minimum non-zero default value	maximum of 10ms of CIR or 6Kbytes

high-prio-only

Syntax **high-prio-only** *percent*
no high-prio-only

Context config>qos>sap-ingress>queue
config>qos>sap-egress>queue

Description The **high-prio-only** command configures the percentage of buffer space for the queue, used exclusively by high priority packets. The specified value overrides the default value for the context.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

The **no** form of this command restores the default high priority reserved size.

Parameters *percent* — The percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10.

Values 0 — 100, default

mbs

Syntax **mbs** *size* [**bytes** | **kilobytes**]
no mbs

Context config>qos>sap-egress>queue
config>qos>sap-ingress>queue

Description The Maximum Burst Size (MBS) command provides the explicit definition of the maximum amount of buffers allowed for a specific queue. The value is given in bytes or kilobytes and overrides the default value for the context.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sap-ingress context for mbs provides a mechanism for overriding the default maximum size for the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS size assigned to the queue to the value.

Default default

Parameters *size* [**bytes** | **kilobytes**] — The size parameter is an integer expression of the maximum number of bytes or kilobytes of buffering allowed for the queue. The default unit is kilobytes; to configure the MBS in bytes specify the bytes parameter. A value of 0 causes the queue to discard all packets.

Values

0 — 131072 or default	
Minimum configurable non-zero value	1byte
Minimum default value	maximum of 10ms of PIR or 64Kbytes

mbs

Syntax **mbs** {[**0..2625**][**kilobytes**] | [**0..2688000**]**bytes** | **default** }
no mbs

Context config>qos>sap-ingress>queue
config>qos>sap-egress>queue

Description The Maximum Burst Size (MBS) command provides the explicit definition of the maximum amount of buffers allowed for a specific queue. The value is given in kilobytes and overrides the default value for the context.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing

packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sap-ingress context for mbs provides a mechanism for overriding the default maximum size for the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The **no** form of this command returns the MBS size assigned to the queue to the value.

Default default

Parameters [0..2625][**kilobytes**] — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps enter the value 100 and specify the **kilobytes** parameter. A value of 0 causes the queue to discard all packets.

Values 0..2625

[0..2688000]**bytes** — The size parameter is an integer expression of the maximum number of bytes or kilobytes of buffering allowed for the queue. For a value of 100 kbps enter the value 100 and specify the **kilobytes** parameter. A value of 0 causes the queue to discard all packets.

Values 0..2688000

default —

packet-byte-offset

Syntax **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}
no packet-byte-offset

Context config>qos>sap-egress>queue

Description This command is used to modify the size of each packet handled by the queue by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed.

The packet-byte-offset command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers.

When a packet-byte-offset value is applied to a queue instance, it adjusts the immediate packet size. This means that the queue rates, i.e., operational PIR and CIR, and queue bucket updates use the adjusted packet size. In addition, the queue statistics will also reflect the adjusted packet size. Scheduler policy rates, which are data rates, will use the adjusted packet size.

The port scheduler max-rate and the priority level rates and weights, if a Weighted Scheduler Group is used, are always on-the-wire rates and thus use the actual frame size. The same goes for the agg-rate-limit on a

SAP, a subscriber, or a Multi-Service Site (MSS) when the queue is port-parented.

When the user enables frame-based-accounting in a scheduler policy or queue-frame-based-accounting with agg-rate-limit in a port scheduler policy, the queue rate will be capped to a user configured on-the-wire rate but the packet-byte-offset value is still in effect as explained above.

The **no** form of this command is used to remove per packet size modifications from the queue.

Parameters **add** *bytes* **—** The **add** keyword is mutually exclusive to the **subtract** keyword. Either parameter must be specified. When **add** is defined, the corresponding bytes parameter specifies the number of bytes that is added to the size of each packet associated with the queue for scheduling and accounting purposes.

Values 0— 32

Default None

subtract *bytes* **—** The **subtract** keyword is mutually exclusive to the **add** keyword. Either parameter must be specified. When subtract is defined, the corresponding bytes parameter specifies the number of bytes that is subtracted to the size of each packet associated with the queue for scheduling and accounting purposes.

Values 0 — 64

Default None

parent

Syntax **parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
no parent

Context config>qos>sap-ingress>queue
config>qos>sap-egress>queue

Description This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the **config>qos>scheduler-policy>tier** *level* context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly or indirectly applied (through a multi-service customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The orphaned state must generate a log entry and a trap message. The SAP which the queue belongs to must also depict an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state mentioned above and automatically return to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of the command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters *scheduler-name* — The defined *scheduler-name* conforms to the same input criteria as the schedulers defined within a scheduler policy. Scheduler names are configured in the `config>qos>scheduler-policy>tier level` context. There are no checks performed at the time of definition to ensure that the *scheduler-name* exists within an existing scheduler policy. For the queue to use the defined *scheduler-name*, the scheduler exists on each egress SAP the queue is eventually created on. For the duration where *scheduler-name* does not exist on the egress SAP, the queue operates in an orphaned state.

Values Any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Default None. Each parental association must be explicitly defined.

weight *weight* — *weight* defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as weighted receive no parental bandwidth until all strict queues and schedulers on the parent have reached their maximum bandwidth or are idle. In this manner, weighted children are considered to be the lowest priority.

All **weight** values from all weighted active queues and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the queue or scheduler after the strict children are serviced. A weight is considered to be active when the pertaining queue or scheduler has not reached its maximum rate and still has packets to transmit. All child queues and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all strict and non-zero weighted queues and schedulers are operating at the maximum bandwidth or are idle.

Values 0 — 100

Default 1

level *level* — The optional **level** parameter defines the level of hierarchy when compared to other schedulers and queues when vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as **strict** receive no parental bandwidth until all strict queues and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

Children of the parent scheduler with a lower strict priority or that are weighted will not receive bandwidth until all children with a higher strict priority have either reached their maximum bandwidth or are idle. Children with the same strict level are serviced in a round robin fashion.

Values 1 — 100

Default 1

cir-weight *cir-weight* **—** Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with

100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 100

cir-level *cir-level* **—** Defines the port priority the queue or scheduler will use to receive bandwidth for its within-cir offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 8 (8 is the highest priority)

Default 0

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*] [**port-limit**|**local-limit**]
percent-rate *pir-percent* **police** [**port-limit**|**local-limit**]
no percent-rate

Context config>qos>sap-egress>queue
config>qos>sap-ingress>queue

Description The percent-rate command within the SAP ingress and egress QOS policy enables supports for a queue's PIR and CIR rate to be configured as a percentage of the egress port's line rate or of its parent scheduler's rate. When the rates are expressed as a port-limit, the actual rates used per instance of the queue will vary based on the port speed. For example, when the same QOS policy is used on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same QOS policy to be used on SAPs on different ports without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

When the rates are expressed as a local-limit, the actual rates used per instance of the queue are relative to the queue's parent scheduler rate. This enables the same QOS policy to be used on SAPs with different parent scheduler rates without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue. If the parent scheduler rate changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the

percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

Queue rate overrides can only be specified in the form as configured in the QoS policy (a SAP override can only be specified as a percent-rate if the associated QoS policy was also defined as percent-rate). Likewise, a SAP override can only be specified as a rate (kbps) if the associated QoS policy was also defined as a rate. Queue-overrides are relative to the limit type specified in the QOS policy.

The **no** form of this command returns the queue to its default shaping rate and cir rate. When no percent-rate is defined within a SAP ingress or egress queue-override, the queue reverts to the defined shaping and CIR rates within the SAP ingress and egress QOS policy associated with the queue.

Parameters *pir-percent* — The pir-percent parameter is used to express the queue's PIR as apercentage dependant on the use of the port-limit or local-limit.

Values Percentage ranging from 0.01 to 100.00. The default is 100.00.

cir *cir-percent* **—** The **cir** keyword is optional and when defined the required cir-percent CIR parameter expresses the queue's CIR as a percentage dependant on the use of the port-limit or local-limit.

Values Percentage ranging from 0.00 to 100.00. The default is 100.00

port-limit — The port**-limit** keyword specifies that the configure PIR and CIR percentages are relative to the rate of the port (including the **ingress-rate**/**egress-rate** setting) to which this queue connects.

local-limit — The local-limit keyword specifies that the configure PIR and CIR percentages are relative to the rate of the queue's parent scheduler **rate** or **agg-rate** rate at egress.

port-parent

Syntax **port-parent** [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
no port-parent

Context config>qos>sap-egress>queue

Description This command specifies whether this queue feeds off a port-level scheduler. When configured, this SAP egress queue is parented by a port-level scheduler. This object is mutually exclusive with SAP egress queue parent. Only one kind of parent is allowed.

The **port-parent** command defines a child/parent association between an egress queue and a port based scheduler or between an intermediate service scheduler and a port based scheduler. The command may be issued in three distinct contexts; **sap-egress queue** *queue-id*, **network-queue queue** *queue-id* and **scheduler-policy scheduler** *scheduler-name*. The **port-parent** command allows for a set of within-cir and above-CIR parameters that define the port priority levels and weights for the queue or scheduler. If the **port-parent** command is executed without any parameters, the default parameters are assumed.

In this context, the **port-parent** command is mutually exclusive to the **parent** command (used to create a parent/child association between a queue and an intermediate scheduler). Executing a **port-parent** command when a parent definition is in place causes the current intermediate scheduler association to be removed and replaced by the defined port-parent association. Executing a **parent** command when a port-parent definition exists causes the port scheduler association to be removed and replaced by the defined intermediate scheduler name.

Changing the parent context on a SAP egress policy queue may cause a SAP or subscriber context of the queue (policy associated with a SAP or subscriber profile) to enter an orphaned state. If an instance of a queue is created on a port or channel that does not have a port scheduler enabled and the sap-egress policy creating the queue has a port-parent association, the queue will be allowed to run according to its own rate parameters and will not be controlled by a virtual scheduling context. If an instance of a queue is on a port or channel that has a port scheduler configured and the sap-egress policy defines the queue as having a non-existent intermediate scheduler parent, the queue will be treated as an orphan and will be handled according to the current orphan behavior on the port scheduler.

The **no** form of this command removes a port scheduler parent association for the queue or scheduler. If a port scheduler is defined on the port which the queue or scheduler instance exists, the queue or scheduler will become orphaned if an port scheduler is configured on the egress port of the queue or scheduler.

Default **no port-parent**

Parameters **weight *weight* —** Defines the weight the queue or scheduler will use at the above-cir port priority level (defined by the level parameter).

Values 0 — 100

Default 1

level *level* — Defines the port priority the queue or scheduler will use to receive bandwidth for its above-cir offered-load.

Values 1 — 8 (8 is the highest priority)

Default 1

cir-weight *cir-weight* — Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 100

cir-level *cir-level* — Defines the port priority the queue or scheduler will use to receive bandwidth for its within-cir offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 8 (8 is the highest priority)

Default 0

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate* | **police**]
no rate

Context config>qos>sap-ingress>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. For SAP ingress, the CIR also defines the rate that packets are considered in-profile by the system. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP ingress or SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default **rate max cir 0 —** The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 — 100000000, **max**

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer.

Values 0 — 100000000, **max**

Default 0

police — Specifies that the out of profile traffic feeding into the physical queue instance should be dropped. Using this keyword will override the bandwidth specified by the SAP ingress queue's administrative CIR.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>qos>sap-egress>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default **rate max cir 0 —** The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been

executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 — 100000000, **max**

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer.

Values 0 — 100000000, **max**

Default 0

xp-specific

Syntax **xp-specific**

Context config>qos>sap-egress>queue

Description This command enables the context to configure IOM3-XP specific information. The xp-specific CLI node within the SAP egress QoS policy queue context is used to specify queue parameters or behavior specific to the Q2 traffic management feature set. All IOMs within the XP family utilize the Q2 for traffic management queuing functions. When the SAP egress QoS policy is applied to a SAP on an IOM3-XP, any commands and parameters defined within the xp-specific context will either override or augment the generic commands and parameters defined for the specific queue ID.

In the event that the QoS policy is applied to a SAP on a non-IOM3-XP, the commands and parameters within the xp-specific node are ignored.

When the QoS policy is applied to a LAG SAP that spans XP and non-XP IOMs, the xp-specific commands and parameters are applied for the SAP queues created on the IOM3-XP LAG links.

wred-queue

Syntax **wred-queue** [**policy** *slope-policy-name*]
no wred-queue

Context config>qos>sap-egress>queue>xp-specific

Description This command alters the generic buffer pool association of the queue for the purpose of allowing queue-specific WRED slopes with minimal provisioning. When the **wred-queue** command is defined and the queue ID is created on an IOM3-XP, a buffer pool is created specifically for the queue and the queue obtains all buffers from that pool. The size of the pool is the same as the size of the queue. In this manner, the WRED slopes that operate based on the pool's buffer utilization are also reacting to the congestion depth of

the queue.

The size of the buffer pool is dictated by the queue's mbs parameter. The size of the reserved CBS portion of the buffer pool is dictated by the queue's cbs parameter. The provisioning characteristics of the **mbs** and **cbs** commands have not been changed.

In the case where the QoS policy is applied to a SAP on an IOM3-XP which has WRED queue support shutdown (**config>card>>fp>egress>wred-queue-control>shutdown**) the queue will continue to map to either to its default pool or the pool defined in the **pool** command. If the **no shutdown** command is executed on the IOM, the queue will at that point be automatically moved to its own WRED pool.

Each pool created for a queue using the **wred-queue** command shares buffers with all other wred-queue enabled queues on the same IOM3-XP. The WRED pool buffer management behavior is defined within the **config>card>fp>egress>wred-queue-control** CLI context.

The WRED slopes within the pool are defined by the slope policy associated with the queue. When a policy is not explicitly defined, the default slope policy is used. The slope policy enables, disables and defines the relative geometry of the high and low WRED slopes in the pool. The policy also specifies the time average factor used by the pool when calculating the weighted average pool depth.

As packets attempt to enter the egress queue, they are associated with either the high or low WRED slope based on the packets profile. If the packet is in-profile, the high slope is used. The low slope is used by out-of-profile packets. Each WRED slope performs a probability discard based on the current weighted average pool depth.

When wred-queue is enabled for a SAP egress queue on an IOM3-XP, the queue pool and hi-priority-only commands are ignored.

The number of wred-queue enabled queues allowed per IOM3-XP is hard coded to 7500.

The **no** form of the command restores the generic buffer pool behavior to the queue. The WRED pool is removed from the system. The queue will be moved to either the default buffer pool or to a named pool if defined and the pool exists.

Default no wred-queue

Parameters **policy** *slope-policy-name* **—** Specifies an existing slope policy that is used to override the default WRED slope policy.

Show Commands

sap-ingress

Syntax **sap-ingress** [*policy-id*] [**detail**]

Context show>qos

Description This command displays SAP ingress QoS policy information.

Parameters *policy-id* — Displays information about the specific policy ID.

Default all SAP ingress policies

Values 1 — 65535

detail — Displays detailed policy information including policy associations.

Sample Output

Show SAP Ingress Output — The following table describes SAP ingress show command output.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Scope	Exclusive – Implies that this policy can only be applied to a single SAP. Template – Implies that this policy can be applied to multiple SAPs on the router.
Description	A text string that helps identify the policy's context in the configuration file.
Default FC	Specifies the default forwarding class for the policy.
Priority	Specifies the enqueuing priority when a packet is marked with a *dot1p-value* specified.
Criteria-type	IP – Specifies that an IP criteria-based SAP ingress policy is used to select the appropriate ingress queue and corresponding forwarding class for matched traffic. MAC – Specifies that a MAC criteria-based SAP is used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.
Mode	Specifies the configured mode of the meter (trTcm or srTcm).

Label	Description (Continued)
CIR Admin	Specifies the administrative Committed Information Rate (CIR) parameters for the queue. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.
CIR Oper	The operational value derived by computing the CIR value from the administrative CIR and PIR values and their corresponding adaptation rules.
CIR Rule	min – The operational CIR for the queue will be equal to or greater than the administrative rate specified using the rate command. max – The operational CIR for the queue will be equal to or less than the administrative rate specified using the rate command. closest – The operational PIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational PIR.
PIR Admin	Specifies the administrative Peak Information Rate (PIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues) or out an egress interface (for SAP egress queues).
PIR Oper	The administrative PIR specified by the user.
PIR Rule	min – The operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command. max – The operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command. closest – The operational PIR for the queue will be the rate closest to the rate specified using the rate command.
CBS	def – Specifies the default CBS value for the queue. value – Specifies the value to override the default reserved buffers for the queue.
MBS	def – Specifies the default MBS value.

Label	Description (Continued)
	`value` – Specifies the value to override the default MBS for the queue.
`HiPrio`	Specifies the percentage of buffer space for the queue, used exclusively by high priority packets.
`PIR Lvl/Wt`	Specifies the priority level of the scheduler when compared to other child schedulers and queue vying for bandwidth on the parent schedulers during the 'above CIR' distribution phase of bandwidth allocation. Weight defines the relative weight of this scheduler in comparison to other child schedulers and queue at the same level.
`CIR Lvl/Wt`	Specifies the level of hierarchy when compared to other schedulers and queue when vying for bandwidth on the parent scheduler. Weight defines the relative weight of this queue in comparison to other child schedulers and queue while vying for bandwidth on the parent scheduler.
`Parent`	Specifies the parent scheduler that governs the available bandwidth given the queue aside from the queue's PIR setting.
`Dot1p`	Specifies the forwarding class or enqueuing priority when a packet is marked with a *dot1p-value* specified.
`FC`	Specifies the forwarding class overrides.
`Priority`	The optional priority setting overrides the default enqueuing priority for the packets received on an ingress SAP which uses the policy that matches this rule. `High` – Specifies that the `high` enqueuing parameter for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. `Low` – Specifies that the `low` enqueuing parameter for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested.
`DSCP`	Specifies the forwarding class or enqueuing priority when a packet is marked with the DiffServ Code Point (DSCP) value.
`FC`	Specifies one of the predefined forwarding classes in the system. When a packet matches the rule the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule.

Label	Description (Continued)
Priority	This parameter specifies the default enqueuing priority overrides for all packets received on an ingress SAP using this policy that match this rule. High – Specifies that the high enqueuing parameter for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. Low – Specifies that the low enqueuing parameter for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested.
Prec	Specifies the forwarding class or enqueuing priority when a packet is marked with an IP precedence value (*ip-prec-value)*.
UCastQ	Specifies the default unicast forwarding type queue mapping.
MCastQ	Specifies the overrides for the default multicast forwarding type queue mapping.
BCastQ	Specifies the default broadcast forwarding type queue mapping.
UnknownQ	Specifies the default unknown unicast forwarding type queue mapping.
Match Criteria	Specifies an IP or MAC criteria entry for the policy.
Entry	
Source IP	Specifies a source IP address range used for an ingress SAP QoS policy match.
Source Port	Specifies a source TCP or UDP port number or port range used for an ingress SAP QoS policy match.
Protocol	Specifies the IP protocol number to be used for an ingress SAP QoS policy match.
DSCP	Specifies a DiffServ Code Point (DSCP) name used for an ingress SAP QoS policy match.
Fragment	True – Configures a match on all fragmented IP packets. False – Configures a match on all non-fragmented IP packets.
FC	Specifies the entry's forwarding class.
Priority	Specifies the default enqueuing priority overrides for all packets received on an ingress SAP using this policy.

Label	Description (Continued)
Src MAC	Specifies a source MAC address or range to be used as a Service Ingress QoS policy match.
Dst MAC	Specifies a destination MAC address or range to be used as a Service Ingress QoS policy match.
Dot1p	Specifies a IEEE 802.1p value to be used as the match.
Snap-pid	Specifies an IEEE 802.3 LLC SNAP Ethernet Frame PID value to be used as a Service Ingress QoS policy match.
Ethernet-type	Specifies an Ethernet type II Ethertype value to be used as a Service Ingress QoS policy match.
ESnap-oui-zero	Specifies an IEEE 802.3 LLC SNAP Ethernet Frame OUI zero or non-zero value to be used as a Service Ingress QoS policy match.
DSAP	Specifies an Ethernet 802.2 LLC DSAP value or range for an ingress SAP QoS policy match.
SSAP	Specifies an Ethernet 802.2 LLC DSAP value or range for an ingress SAP QoS policy match.
FC	Specifies the entry's forwarding class.
Priority	Specifies the default enqueuing priority overrides for all packets received on an ingress SAP using this policy.
Service Association	
Service-Id	The unique service ID number which identifies the service in the service domain.
Customer-Id	Specifies the customer ID which identifies the customer to the service.
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP ingress policy is applied.

```
QoS Sap Ingress
===============================================================================
Sap Ingress Policy (100)
-------------------------------------------------------------------------------
Policy-id      : 100                          Scope        : Template
Default FC     : be                           Priority     : Low
Criteria-type  : IP
Description    : Used on VPN sap
-------------------------------------------------------------------------------
Queue Mode    CIR Admin PIR Admin CBS     HiPrio  PIR Lvl/Wt    Parent
              CIR Rule  PIR Rule  MBS             CIR Lvl/Wt
```

```
-------------------------------------------------------------------------------
1     Prio      0         max     def    def           1/1           None
                closest   closest def                  0/1
2     Prio      0         max     def    def           1/1           None
                closest   closest def                  0/1
10    Prio      0         11000   def    def           1/1           VPN_be
                closest   closest def                  0/1
11    Prio      0         max     def    def           1/1           None
                closest   closest def                  0/1
12    Prio      0         11000   def    def           1/1           VPN_prio*
                closest   closest def                  0/1
13    Prio      0         1       def    def           1/1           VPN_rese*
                closest   closest def                  0/1
15    Prio      1500      1500    def    def           1/1           VPN_video
                closest   closest def                  0/1
16    Prio      2500      2500    def    def           1/1           VPN_voice
                closest   closest def                  0/1
17    Prio      36        100     def    def           1/1           VPN_nc
                closest   closest def                  0/1
20    Prio      0         11000   def    def           1/1           VPN_be
                closest   closest def                  0/1
22    Prio      0         11000   def    def           1/1           VPN_prio*
                closest   closest def                  0/1
23    Prio      0         1       def    def           1/1           VPN_rese*
                closest   closest def                  0/1
25    Prio      1500      1500    def    def           1/1           VPN_video
                closest   closest def                  0/1
26    Prio      2500      2500    def    def           1/1           VPN_voice
                closest   closest def                  0/1
27    Prio      36        100     def    def           1/1           VPN_nc
                closest   closest def                  0/1
-------------------------------------------------------------------------------
FC                  UCastQ          MCastQ          BCastQ          UnknownQ
-------------------------------------------------------------------------------
be                  10              20              20              20
af                  12              22              22              22
h2                  16              26              26              26
ef                  13              23              23              23
h1                  15              25              25              25
nc                  17              27              27              27
-------------------------------------------------------------------------------
SubFC                               Profile         In-Remark       Out-Remark
-------------------------------------------------------------------------------
af                                  None            None            None
be                                  None            None            None
ef                                  None            None            None
h1                                  None            None            None
h2                                  None            None            None
nc                                  None            None            None
-------------------------------------------------------------------------------
Dot1p           FC                              Priority
-------------------------------------------------------------------------------
0               af                              High
1               ef                              High
7               be                              Low
-------------------------------------------------------------------------------
DSCP            FC                              Priority
-------------------------------------------------------------------------------
af41            af                              High
```

```
-------------------------------------------------------------------------------
Prec Value      FC                                  Priority
-------------------------------------------------------------------------------
0               be                                  Default
2               af                                  Default
3               ef                                  Default
5               h1                                  Default
6               h2                                  Default
7               nc                                  Default
-------------------------------------------------------------------------------
Match Criteria
-------------------------------------------------------------------------------
IP Match Criteria
-------------------------------------------------------------------------------
Entry          : 10
Description    : Entry 10-FC-AF
Source IP      : 10.10.10.104/24             Source Port  : None
Dest. IP       : Undefined                   Dest. Port   : None
Protocol       : 6                           DSCP         : None
Fragment       : Off
FC             : af                          Priority     : High

Entry          : 20
Description    : Entry 20-FC-BE
Source IP      : Undefined                   Source Port  : None
Dest. IP       : Undefined                   Dest. Port   : eq 255
Protocol       : 17                          DSCP         : None
Fragment       : Off
FC             : Default                     Priority     : Default
-------------------------------------------------------------------------------
IPv6 Match Criteria
-------------------------------------------------------------------------------
No Match Criteria Entries found.
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 700 (VPLS)                  Customer-Id  : 7
 - SAP : 1/1/9:0                  override
===============================================================================
*A:ALA-48>config>qos#


config>qos# show qos sap-ingress 2 detail
===============================================================================
QoS Sap Ingress
-------------------------------------------------------------------------------
Sap Ingress Policy (2)
-------------------------------------------------------------------------------
Policy-id      : 2                           Scope        : Template
Default FC     : be                          Priority     : Low
Criteria-type  : None
-------------------------------------------------------------------------------
Queue Mode    CIR Admin PIR Admin CBS     HiPrio  PIR Lvl/Wt    Parent
              CIR Rule  PIR Rule  MBS             CIR Lvl/Wt
-------------------------------------------------------------------------------
1     Prio      0         max     def    def           1/1           None
                closest   closest def                  0/1
11    Prio      0         max     def    def           1/1           None
                closest   closest def                  0/1
```

```
-------------------------------------------------------------------------
FC                  UCastQ         MCastQ         BCastQ         UnknownQ
-------------------------------------------------------------------------
af                  def            def            def            def
ef                  def            def            def            def
-------------------------------------------------------------------------
SubFC         DE-1-out-profile     Profile        In-Remark      Out-Remark
-------------------------------------------------------------------------
af            No                   None           None           None
ef            Yes                  None           None           None
-------------------------------------------------------------------------
Dot1p           FC                              Priority
-------------------------------------------------------------------------
No Dot1p-Map Entries Found.
-------------------------------------------------------------------------
DSCP            FC                              Priority
-------------------------------------------------------------------------
No DSCP-Map Entries Found.
-------------------------------------------------------------------------
Prec Value      FC                              Priority
-------------------------------------------------------------------------
No Prec-Map Entries Found.
-------------------------------------------------------------------------
Match Criteria
-------------------------------------------------------------------------
No Matching Criteria.
-------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------
No Associations Found.
config>qos#

# show qos sap-ingress

===============================================================================
Sap Ingress Policies
===============================================================================
Policy-Id  Scope     Name                     Description
-------------------------------------------------------------------------------
1          Template  default                  Default SAP ingress QoS policy.
10         Template
20         Template
-------------------------------------------------------------------------------
Number of policies : 3
-------------------------------------------------------------------------------
===============================================================================
*A:#
```

sap-egress

Syntax **sap-egress** [*policy-id*] [**detail**]

Context show>qos

Description This command displays SAP egress QoS policy information.

Parameters *policy-id* — Displays information about the specific policy ID.

Values 1 — 65535

detail — Displays detailed policy information including policy associations.

SAP Egress Output — The following table describes SAP egress show command output.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Scope	Exclusive – Implies that this policy can only be applied to a single SAP. Template – Implies that this policy can be applied to multiple SAPs on the router.
Description	A text string that helps identify the policy's context in the configuration file.
Queue:	
CIR Admin	Specifies the administrative Committed Information Rate (CIR) parameters for the queue. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.
CIR Oper	The operational value derived by computing the CIR value from the administrative CIR and PIR values and their corresponding adaptation rules.
CIR Rule	min – The operational CIR for the queue will be equal to or greater than the administrative rate specified using the rate command except where the derived operational CIR is greater than the operational PIR. If the derived operational CIR is greater than the derived operational PIR, the operational CIR will be made equal to the operational PIR. max – The operational CIR for the queue will be equal to or less than the administrative rate specified using the rate command.

Label	Description (Continued)
	`closest` – The operational PIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational PIR.
`PIR Admin`	Specifies the administrative Peak Information Rate (PIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues) or out an egress interface (for SAP egress queues).
`PIR Oper`	The administrative PIR specified by the user.
`PIR Rule`	`min` – The operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command. `max` – The operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command. `closest` – The operational PIR for the queue will be the rate closest to the rate specified using the rate command.
`CBS`	`def` – Specifies that the CBS value reserved for the queue. `value` – Specifies the value to override the default reserved buffers for the queue.
`MBS`	`def` – Specifies that the MBS value is set by the def-mbs function. `value` – Specifies the value to override the default maximum size for the queue.
`HiPrio`	Specifies the percentage of buffer space for the queue, used exclusively by high priority packets.
`PIR Lvl/Wt`	Specifies the priority level of the scheduler when compared to other child schedulers and queues vying for bandwidth on the parent schedulers during the 'above CIR' distribution phase of bandwidth allocation. Weight defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same level.
`CIR Lvl/Wt`	Specifies the level of hierarchy when compared to other schedulers and queues when vying for bandwidth on the parent scheduler. Weight defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent scheduler.

Label	Description (Continued)
Parent	Specifies the parent scheduler that governs the available bandwidth given the queue aside from the queue's PIR setting.
FC Name	Specifies the forwarding class queue mapping or dot1p marking is to be edited.
Queue-id	Specifies the *queue-id* that uniquely identifies the queue within the policy.
Explicit/Default	Explicit – Specifies the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. Default – Specifies that the default dot1p value (0) is used.
Service Association	
Service-Id	The unique service ID number which identifies the service in the service domain.
Customer-Id	Specifies the customer ID which identifies the customer to the service.
SAP	Specifies the a Service Access Point (SAP) within the service where the policy is applied.
Mirror SAPs:	
Mirror Dest	Specifies the mirror service ID which identifies the service in the service domain.
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP egress policy is applied.

Sample Output

```
A:ALA-49# show qos sap-egress
===============================================================================
Sap Egress Policies
===============================================================================
Policy-Id            Scope     Description
-------------------------------------------------------------------------------
1                    Template  Default SAP egress QoS policy.
1010                 Template
1020                 Template
===============================================================================
A:ALA-49#


A:ALA-49# show qos sap-egress 1010
===============================================================================
QoS Sap Egress
===============================================================================
```

```
-------------------------------------------------------------------------------
Sap Scheduler Policy (1010)
-------------------------------------------------------------------------------
Policy-id      : 1010                          Scope        : Template
===============================================================================
A:ALA-49#


A:ALA-49# show qos sap-egress 1010 detail
===============================================================================
QoS Sap Egress
===============================================================================
-------------------------------------------------------------------------------
Sap Scheduler Policy (1010)
-------------------------------------------------------------------------------
Policy-id      : 1010                          Scope        : Template
-------------------------------------------------------------------------------
Queue          CIR Admin PIR Admin CBS     HiPrio  PIR Lvl/Wt    Parent
               CIR Rule  PIR Rule  MBS             CIR Lvl/Wt
-------------------------------------------------------------------------------
1              0         max       def     def     1/1           None
               closest   closest   def             0/1
8              0         max       def     def     1/1           None
               closest   closest   def             0/1
-------------------------------------------------------------------------------
FC Name               Queue-id    Explicit/Default
-------------------------------------------------------------------------------
be                    8           Explicit (7)
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 1 (VPRN)                      Customer-Id  : 1
 - SAP : 1/1/10:1

SLA Profiles :
 - test                             override
-------------------------------------------------------------------------------
Mirror SAPs
-------------------------------------------------------------------------------
No Mirror SAPs Found.
===============================================================================
A:ALA-49#

config>qos# show qos sap-egress 2 detail
===============================================================================
QoS Sap Egress
-------------------------------------------------------------------------------
Sap Scheduler Policy (2)
-------------------------------------------------------------------------------
Policy-id      : 2                             Scope        : Template
-------------------------------------------------------------------------------
Queue CIR Admin PIR Admin CBS     HiPrio PIR Lvl/Wt    Parent         AvgOvrhd
      CIR Rule  PIR Rule  MBS             CIR Lvl/Wt
-------------------------------------------------------------------------------
1     0         max       def     def    1/1           None           0.00
      closest   closest   def            0/1
-------------------------------------------------------------------------------
FC Name               Queue-id    Explicit/Default         DE-Mark
-------------------------------------------------------------------------------
```

```
af                    def         Explicit (4)              Profile
l1                    def         Explicit (In:5 Out:6)     Force 0
ef                    def         Default                   None
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
No Associations Found.
-------------------------------------------------------------------------------
Mirror SAPs
-------------------------------------------------------------------------------
No Mirror SAPs Found.
===============================================================================
config>qos#
```

queue

Syntax **queue from** {**sap** *sap-id* | **queue-group** *port-id queue-group-name* | **subscriber** *subscriber-id* | **network** {*mda-id* | *port-id*} | **system** {**card** *slot-number* | **mda** *mda-id* **port** *port-id*}} {**ingress** | **egress**} [**id** *queue-id*]

Context show>qos

Description The show qos queue command outputs the Burst Control Group (BCG) name and slowest accurate visitation time for the specified queues.

For each queue specified, the system may find multiple hardware queues. This may be true for ingress queues on multipoint services (VPLS, IES, VPRN) or for queues created on an Ethernet Link Aggregation Group (LAG). When this is true, the show command may display the calculated slowest accurate visitation time for the logical queue (all hardware queues will have the same calculated value) but must display the BCG name for each individual hardware queue.

The BCG name associated with a queue may be specified in the show bcg command to display the historical and current visitation time for the BCG managing the burst tolerance of the queue. If the output visitation time is greater (longer time) than the queue returned slowest accurate visitation time, the queue's shaping rate may be negatively impacted.

Parameters **from —** The from keyword specifies that the following parameters are match criteria for finding a single or set of ingress or egress queues within the system. The system will accept sap, queue-group, subscriber, network-queues or shared-queues as the match criteria.

sap *sap-id* **—** The sap keyword is used to specify that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues within the specified sap-id. The sap keyword is mutually exclusive with the other from match criteria. If the specified sap-id is not found, the system should return 'The specified SAP ID does not exist'.

queue-group *port-id queue-group-name* **—** The queue-group keyword is used to specify that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues within the specified queue-group-name on the specified port-id. The following ingress or egress keyword further specifies that the targeted queue group is an ingress port or egress port queue group. The queue-group keyword is mutually exclusive with the other from match criteria. If the specified port-id is not provisioned on the system or the specified queue-group-name is not found on the ports specified direction, the system should return 'The specified queue group does not exist'.

subscriber *subscriber-id* **—** The subscriber keyword is used to specify that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues associated with the specified subscriber-id. The queue-group keyword is mutually exclusive with the other from match criteria. If the specified subscriber-id does not exist, the system should return 'The specified subscriber does not exist'.

network {*mda-id* **|** *port-id*} **—** The network keyword is used to specify that the system should find and display the queue informationfor the queues associated with the specified mda-id or port-id. If the ingress direction qualifier is specified, an mda-id is required. If the egress direction qualifier is specified, a port-id is required. The network keyword is mutually exclusive with the other from match criteria. If the specified mda-id does not exist, the system should return 'The specified MDA is not provisioned'. If the specified port-id does not exist, the system should return 'The specified port is not provisioned'.

system {**card** *slot-number* |*mda-id* **|** *port-id*} **—** The system keyword is used to specify that the system should find and display the queue information for all the system queues associated with the specified card slot-id, mda mda-id or port port-id. If the ingress direction qualifier is specified, the ingress system queues are displayed. If the egress direction qualifier is specified, only the egress system queues are displayed. The system keyword is mutually exclusive with the other from match criteria. If the specified slot-id does not exist, the system should return 'The specified slot number is not provisioned'. If the specified mda-id does not exist, the system should return 'The specified MDA is not provisioned'. If the specified port-id does not exist, the system should return 'The specified port is not provisioned'. The id parameter is not supported when matching system queues.

{**ingress | egress**} **—** The ingress and egress direction qualifiers are mutually exclusive. Either ingress or egress must be specified.

id queue-id — The id keyword is used to limit the return queues to a single queue-id. The keyword is not accepted when the system match criteria is used.

bcg

Syntax **bcg** *burst-control-group-name* [**member-queues** [**at-risk-only**]] [**exp-util-bw** *megabits-per-second*]

Context show>qos

Description The show qos bcg command outputs the current and historical visitation time associated with the specified BCG name.

A Burst Control Group (BCG) represents a list of queues that share the same non-scheduling PIR and CIR bucket target update interval. When a queue's scheduled rate bursts above its PIR bucket depth, the queue is removed from its scheduling context. The system uses a BCG in order to visit the queues PIR bucket to periodically drain an appropriate amount from the bucket. When the bucket has been drained below the PIR bucket threshold, the queue is allowed back onto its scheduling context. The amount decremented from the bucket is a function of the amount of time that has elapsed since the last bucket update and the queue's shaping rate (PIR). If the queue's shaping rate is configured as 1Mbps and 1ms has elapsed since the last bucket update, the system will decrement the PIR bucket by 125 bytes. One caveat is that the bucket cannot be decremented past a depth of 0. This fact drives how the system chooses which BCG is used to manage the queue bucket update interval.

If a queue's shaping rate is 1Mbps and the threshold (burst limit) is set to 10Kbytes, the maximum amount of time that can expire before the queue is updated without resulting in a negative bucket depth is 81.92ms. This can be calculated by taking the number of bits represented by the bucket depth (10Kbytes = 10 * 1,024 * 8 = 81,920 bits) and dividing it by the rate (81,920 bits / 1,000,000 bits per second = 81.92ms). The queue will not be removed from the scheduler until the PIR bucket depth has equaled or exceeded the configured burst threshold, so the bucket will be at least 10Kbytes deep. If the system visits the queue PIR bucket within 81.92ms, the resulting decrement operation will leave the bucket. If the system takes longer than 81.92ms, the decrement result will be greater than 10Kbytes and part of the decrement result will be lost. The net result is from less than timely updates is that the queue will not be returned to the scheduler context fast enough and some shaping bandwidth for the queue will be lost (underrun the shaping rate).

Each Q2 based forwarding plane maintains 7 Burst Control Groups, each targeting a certain queue bucket visitation time. A 40ms, 20ms, 10ms, 5ms, 1ms, 500us and 100us BCG is supported. By default, queues are placed on a BCG based on shaping rate and the queue's burst limit (PIR threshold depth) is set based on the BCG visitation time and the queue's specified shaping rate. When all shaping queues on a Q2 are left in a default burst tolerance management state, the system has sufficient BCG visitation resources to ensure that all queues do not experience inaccurate bucket decrement conditions.

When explicit burst-limit threshold values are defined for a shaping queue, the system picks an appropriate BCG based on the queue's configured shaping rate and the explicit threshold to find a BCG with the best target visitation time that results in worst case decrement values that are less than the configured threshold. However, when a queue is placed on a 'faster' BCG, more visitation resources are consumed and it is possible that the system will not meet a queue's decrement constraints.

The **show qos bcg** command allows visibility into a BCG's historic and current visitation time. The system samples the amount of time it takes each list to visit each of its associated queues once each second and stores the last 10 samples. It also keeps the longest visitation time seen since the last time the BCG statistics were cleared, the longest visitation time for the current queue-to-BCG lists associations, calculated longest visitation time based on maximum scheduling bandwidth and lastly the longest visitation time for an optionally defined scheduling rate.

With each sample, the system indirectly calculates the amount of scheduling bandwidth based on how much Q2 resources were diverted from BNG visitation processing. This calculated scheduling bandwidth is useful since it can be used to evaluate the worst case longest visitation times for each BCG. The calculated scheduling bandwidth value is stored with the longest seen visitation time and the longest seen visitation time with the current queue-to-BCG mappings.

Parameters *burst-control-group-name* — The burst-control-group-name is required and specifies which globally unique Burst Control Group will be displayed. If the specified Burst Control Group does not exist, the show command will fail and the system will return 'The specified BCG does not exist'.

member-queues [**at-risk-only**] **—** The member-queues optional keyword is used to include a list of all queues attached to the specified burst-control-group-name. The optional at-risk-only keyword may be added to limit the displayed queues to only include queues that are considered 'at-risk' for inaccurate shaping based on either the 100% worst case scheduling bandwidth for the current queue mappings. The 100% scheduling bandwidth used in the 'at-risk' determination may be overridden with a specified scheduling bandwidth by using the exp-util-bw parameter.

exp-util-bw — An optional keyword used to display a calculated worst case visitation rate for the specified burst-control-group-name based on the specified value for megabits-per-second.

megabits-per-second — A value also modifies the member-queues 'at-risk' state output.

Queue Sharing and Redirection

In This Section

This section provides information to configure queue groups using the command line interface.

Topics in this section include:

- Queue Sharing and Redirection on page 380
  - → Pseudowire Shaping for Layer 2 and Layer 3 Services on page 383
- Queue Group Templates and Port Queue Groups on page 385
- Access SAP Forwarding Class Based Redirection on page 389
- Network IP Interface Forwarding Class-Based Redirection on page 397
- Queue Group Behavior on LAG on page 410
- Configuring Ingress/Egress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection on page 427
- Basic Configurations on page 413

Queue Sharing and Redirection

Queue groups are objects created on access or network Ethernet port or ingress forwarding plane of an IOM/IMM/XMA that allow SAP or IP interface forwarding classes to be redirected from the normal type of queue mapping to a shared queue. Queue groups may contain queues, policers, or a combination or the two depending on the type of queue group.The following types of queue groups are supported:

- Access ingress supports a single queue group instance per ingress port, or multiple queue groups created at the ingress forwarding plane level of the IOM/IMM/XMA. Access ingress port queue groups may only contain queues, whereas access ingress forwarding plane queue groups may only contain policers.
- Access egress supports the creation of multiple queue groups per egress port. These queue groups may only contain queues.
- Network ingress supports the creation of multiple queue groups at the ingress forwarding plane level of the IOM/IMM/XMA. These queue groups may only contain policers.
- Network egress supports the creation of multiple queue groups per egress port. These queue groups may contain queues, only, or queues and policers.

Supported Platforms

Queue sharing and redirection is supported on the SR and ESS platforms with the following IOM types:

- Access SAP port queue group supported on IOM-1 of types the iom-10g, iom-20g, and iom- 20g-b. Network queue groups are not supported.
- Access SAP port and network port queue group are supported on IOM-2s. Up to 20K SAPs per MDA can be configured with any supported Ethernet MDA.
- Access SAP port and ingress forwarding plane and network port and ingress forwarding plane queue groups are supported on IOM-3s.

Queue sharing and redirection are also supported in conjunction with the use of existing Ethernet MDA, Ethernet CMA, HS-MDA and the VSM MDA.

Queue Group Applications

Access SAP Queue Group Applications

Normally, each SAP (Service Access Point) has dedicated ingress and egress queues that are only used by that particular SAP. The SAP queues are created based on the queue definitions within the SAP ingress and SAP egress QoS policy applied to the SAP. Each packet that enters or egresses the SAP has an associated forwarding class. The QoS policy is used to map the forwarding class to one of the SAP's local queue IDs. This per-SAP queuing has advantages over a shared queuing model in that it allows each SAP to have a unique scheduling context per queue. During congestion, SAPs operating within their conforming bandwidth will experience little impact since they do not need to compete for queue buffer space with misbehaving or heavily loaded SAPs.

The situation is different for a shared or port-queuing model that is based on policing color packets that conform or exceed a static rate before the single queue and that use WRED or drop tail functions to essentially reserve room for the conforming packets.

In this model, there is no way for the conforming packets to go to the head of line in the view of the port scheduler. Another advantage of per-SAP queuing is the ability for the SAP queues to perform shaping to control burst sizes and forwarding rates based on the SAPs defined SLA. This is especially beneficial when a provider is enforcing a sub-line rate bandwidth limit and the customer does not have the ability to shape at the CE.

However, there are cases where per-SAP queuing is not preferred. Per SAP queuing requires a more complex provisioning model in order to properly configure the SAPs ingress and egress SLAs. This requires service awareness at some points in the network where an aggregation function is being performed. In this case, a shared queuing or per-port queuing model will suffice. Creating ingress and egress access queue groups and mapping the SAPs forwarding classes to the queues within the queue group provides this capability.

A further use case is where a set of ingress SAPs, which may represent a subset of the total number of ingress SAPs, is to be shaped or policed on an aggregate per-forwarding class basis when those SAPs are spread across a LAG on multiple ingress ports, and where color-aware treatment is required so that explicitly in-profile traffic is honored up to CIR, but above which it is marked as out of profile

The above scenarios can be supported with access queue groups. A single ingress queue group is supported per access port, while multiple ingress queue group instances are supported per IOM/IMM/XMA forwarding plane. To provide more flexibility on the egress side of the access port, multiple egress access queue group queue-group instances are supported per egress access port.

Since queue redirection is defined per forwarding class, it is possible to redirect some forwarding classes to a queue group while having others on the SAP use the SAP local queues. This is helpful when shared queuing is only desired for a few applications such as VOIP or VOD while other applications still require queuing at the SAP level.

Network Port Queue Groups for IP Interfaces

Queue groups may be created on egress network ports in order to provide network IP interface queue redirection. A single set of egress port based forwarding class queues are available by default and all IP interfaces on the port share the queues. Creating a network queue group allows one or more IP interfaces to selectively redirect forwarding classes to the group in order to override the default behavior. Using network egress queue groups it is possible to provide dedicated queues for each IP interface.

Note that non-IPv4/non-IPv6/non-MPLS packets will remain on the regular network port queues. Therefore, when using an egress port-scheduler it is important to parent the related regular network port queues to appropriate port-scheduler priority levels to ensure the desired operation under port congestion. This is particularly important for protocol traffic such as LACP, EFM-OAM, ETH-CFM, ARP and IS-IS, which by default use the FC NC regular network port queue.

Pseudowire Shaping for Layer 2 and Layer 3 Services

This feature allows the user to perform ingress and egress data path shaping of packets forwarded within a spoke-sdp (PW). It applies to a VLL service, a VPLS/B-VPLS service, and an IES/VPRN spoke-interface.

The ingress PW rate-limiting feature uses a policer in the queue-group provisioning model. This model allows the mapping of one or more PWs to the same instance of policers that are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress PW shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC that needs to be redirected and optionally for each traffic type (unicast or multicast).
2. Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface that the PW packets can be received on. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different PWs to different queue-group templates.
4. Apply this network QoS policy to the ingress context of a spoke-sdp inside a service, or to the ingress context of a PW template and specify the redirect queue-group name.

One or more spoke-sdps can have their FCs redirected to use policers in the same policer queue-group instance.

The egress PW shaping provisioning model allows the mapping of one or more PWs to the same instance of queues, or policers and queues, that are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only, or policers and queues for each FC which needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface which the PW packets can be forwarded on. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified

in this step, which means the same network QoS policy can redirect different PWs to different queue-group templates.

4. Apply this network QoS policy to the egress context of a spoke-sdp inside a service, or to the egress context of a PW template and specify the redirect queue-group name.

One or more spoke-sdps can have their FCs redirected to use queues only, or queues and policers in the same queue-group instance.

Queue Group Templates and Port Queue Groups

Queue Group Templates

Before a queue group with a specific name may be created on a port or an IOM/IMM/XMA ingress forwarding plane, a queue group template with the same name must first be created. The template is used to define each queue, scheduling attributes and its default parameters. When a queue or policer is defined in a queue group template, that queue will exist in every instance of a port or forwarding plane queue group with that template name. The default queue or policer parameters (such as rate or mbs values) may be overridden with a specific value in each queue group. This works in a similar manner to SAP ingress or SAP egress QoS policies.

Queue sharing is also supported if the High Scale MDA (HSMDA) is used. On ingress, HSMDA queues are bypassed, and the queue group on the IOM forwarding plane is used. On egress, it is possible to redirect forwarding classes from multiple SAPs to an HSMDA queue group. Note that the HSMDA also uses the term *queue group* to describe a group of 8 pre-configured hardware queues on its egress port. When queue sharing and redirection is configured on egress, a set of 8 HSMDA queues could be configured as a part of the queue group template. These correspond to 8 hardware queues on the HSMDA. When all eight (8) egress fcs are mapped to the queue-group instantiated in the egress port, the per-sap hsmda queue-group resource is freed.

Port Queue Groups

Once an ingress or egress queue group template is defined, a port based queue group with the same name may be created. Port queue groups are named objects that act as a container for a group of queues. The queues are created based on the defined queue IDs within the associated queue group template. Port queue groups must be created individually on the ingress and egress sides of the port, but multiple port queue groups of the same template name may be created on egress ports if they have a different instance identifier. These are termed ‘queue group instances’. Each instance of a named queue group created on a port is an independent set of queues structured as per the queue group template. Port queue groups are only supported on Ethernet ports and may be created on ports within a LAG.

Percent-Rate Support

The **percent-rate** command is supported in a queue group template for **pir** and **cir** parameters only for egress queues. The user has the option of specifying **percent-rate** for **pir** and **cir** parameters. For **pir**, the range is 0.01 to 100.00, and for **cir**, the range is 0.00 to 100.00.

The rate can be also configured using the existing keyword rate in Kbps.

When the queue rate is configured with **percent-rate**, a port-limit is applied, specifically, the **percent-rate** is relative to the rate of the port to which the queue is attached.

```
*A:PE>config>qos>qgrps>egr>qgrp>queue# percent-rate
  - no percent-rate
  - percent-rate <pir-percent> [cir <cir-percent>]

 <pir-percent>        : [0.01..100.00]
 <cir-percent>        : [0.00..100.00]
```

Forwarding Plane Queue Groups

Ingress forwarding plane queue groups allow groups of SAPs on one or more ports, or on a LAG on the IOM/IMM/XMA, to be bundled together from a QoS enforcement perspective with an aggregate rate limit to be enforced across all SAPs of a bundle. Multiple queue groups are supported per IOM/IMM/XMA or port on access ingress. These are implemented at the forwarding plane level on the ingress IOM so that SAPs residing on different ingress ports or SAPs on a LAG spread across ports on a given IOM can be redirected to the same queue group

Once an ingress queue group template is defined, a forwarding plane queue group with the same name may be created on an ingress forwarding plane of an IOM/IMM/XMA. Forwarding plane queue groups are named objects that act as a container for a group of policers. Queues are not supported in forwarding plane queue groups. Only hierarchical policers are supported in the forwarding plane queue group, rather than queues. These policers may be configured to use profile-aware behavior. The policers are created based on the defined policer IDs within the associated queue group template. Multiple forwarding plane queue groups of the same template name may be created on ingress if they have a different instance identifier. These are termed *queue group instances*. Each instance of a named queue group created on a forwarding plane is an independent set of policers structured as per the queue group template. Forwarding plane queue groups are only supported with Ethernet ports and may be created on IOM/IMM/XMAs with ports in a LAG.

Redirection Models

Two models are supported for forwarding class redirection. In the first, the actual instance of a queue group to use for forwarding class redirection is named in the QoS policy. This is termed *policy-based redirection*.

In the second model, the forwarding class queue or policers to apply redirection to are identified in the ingress or egress QoS policy. However, the specific named queue group instance is not identified until a QoS policy is applied to a SAP. This is termed *SAP-based redirection.*

Policy-based redirection allows different forwarding classes in the same QoS policy to be redirected to different queue groups, but it requires at least one QoS policy to be configured per queue group instance.

SAP-based redirection can require less QoS policies to be configured since the policy does not have to name the queue group. However, if redirected, all forwarding classes of a given SAP must use the same named queue group instance.

Policy based redirection is applicable to port queue groups on access ingress and access and network egress, while SAP based redirection is applicable to forwarding plane queue groups on access and network ingress, and port queue groups on access and network egress.

Access SAP Forwarding Class Based Redirection

Forwarding class redirection is provisioned within the SAP ingress or SAP egress QoS policy. In each policy, the forwarding class to queue ID mapping may optionally specify a named queue group instance (policy-based redirection) or may simply tag the forwarding class for redirection (SAP-based redirection). When the name is specified, the defined queue ID must exist in the queue group template with the same name.

Redirecting a SAP forwarding class to a queue within a port based queue group using policy-based redirection requires four steps:

1. Create an ingress or egress queue group template. If the forwarding class redirection is in the ingress SAP path, an ingress queue group template must be created. Similarly, an egress queue group template must be created for egress forwarding class redirection. Optionally, you can create the queues in a template by using default parameters. Individual queues must be created before they are associated with a forwarding class. The default queue parameters may be overridden on each port based queue group.

2. Create an ingress or egress queue group instance with the same name as the template on the port associated with the SAP. Examples are as follows:

   On ingress ports:
   config>port>ethernet>access>ingress>queue-group *queue-group-name*

   On egress ports:
   config>port>ethernet>access>egress>queue-group *queue-group-name* [**instance** *instance-id*]

   Queue parameter overrides can also be applied at this time.

3. Redirect the SAP ingress or SAP egress QoS policy forwarding class policer or queue to the queue group name and desired queue ID (Steps 2 and 3 may be done in opposite order). Examples are as follows:

   On ingress:
   config>qos>sap-ingress *policy-id*
   fc *fc-name*
   queue *queue-id* **group** *queue-group-name*

   On egress:
   config>qos>sap-egress *policy-id*
   fc *fc-name*
   queue *queue-id* **group** *queue-group-name* **instance** *instance-id*

config>qos>sap-egress *policy-id*
fc *fc-name*
policer *policer-id* **group** *queue-group-name* **instance** *instance-id*

4. Finally, the SAP ingress or SAP egress QoS policy must be applied to the SAP.

Redirecting a SAP forwarding class to a queue within an egress port based or ingress forwarding plane queue group using SAP-based redirection requires four steps:

1. Create an ingress or egress queue group template. If the forwarding class redirection is in the ingress SAP path, an ingress queue group template must be created. Similarly, an egress queue group template must be created for egress forwarding class redirection. Optionally, you can create the queues in a template by using default parameters. Individual queues must be created before they are associated with a forwarding class. The default queue parameters may be overridden on each port based queue group.

2. Create an ingress queue group instance on the forwarding plane of the IOM/IMM/XMA, or an egress port queue group with the same name as the template on the port associated with the SAP.

   On ingress:
   config>card>fp>ingress>access>queue-group *queue-group-name* **instance** *instance-id* [**create**]

   On egress:
   config>port>ethernet>access>egress>queue-group *queue-group-name* [**instance** *instance-id*]

3. Redirect the SAP ingress forwarding class policer in the SAP-ingress QoS policy using the keyword **fp-redirect-group** keyword on the policer, or SAP egress forwarding class queue or policer using the **port-redirect-group** keyword. (Steps 2 and 3 may be done in opposite order.)

   On ingress:
   config>qos>sap-ingress *policy-id*
   fc *fc-name*
   queue *queue-id* **fp-redirect-group**

   On egress:
   config>qos>sap-egress *policy-id*
   fc *fc-name*
   queue *queue-id* **port-redirect-group-queue**

   config>qos>sap-egress *policy-id*
   fc *fc-name*
   policer *policer-id* **port-redirect-group-queue**

4. Finally, the SAP ingress or SAP egress QoS policy must be applied to the SAP. The named queue group instance that was created on the ingress forwarding plane or the egress port must be specified at this time.

   On ingress:
   config>service>epipe>sap *sap-id*
   ingress
   qos *sap-ingress-policy-id* **fp-redirect-group** *queue-group-name*
   instance *instance-id*

   On egress:
   config>service>epipe>sap *sap-id*
   egress
   qos *sap-egress-policy-id* **port-redirect-group** *queue-group-name*
   instance *instance-id*

Note that redirection to a queue group on the HSMDA supports the SAP-based provisioning model, only.

Ingress and Egress SAP Forwarding Class Redirection Association Rules

Policy Based Provisioning Model

The association rules between SAP ingress and egress QoS policies and queue group templates are simple since both the target queue group name and queue ID within the group are explicitly stated within the access QoS policies.

The following association rules apply when the policy based provisioning model is applied with port queue groups.

When a SAP ingress QoS policy forwarding class is redirected to a queue group queue ID:

- If the queue group name does not exist as an ingress queue group template, the forwarding class redirection will fail.
- If a redirection queue ID does not exist within the ingress queue group template, the forwarding class redirection will fail.
- If the SAP ingress QoS policy is currently applied to a non-Ethernet port or an Ethernet port where the specified ingress queue group does not exist, the forwarding class redirection will fail.

When a SAP ingress QoS policy forwarding class redirection is removed from a queue group queue ID:

- If the forwarding class is being moved to another queue group queue ID that does not exist within an ingress queue group template, the redirection removal from the current queue group queue ID will fail.
- If the forwarding class is being moved to a local queue ID within the SAP ingress QoS policy and the local queue ID does not exist, the redirection removal from the current queue group queue ID will fail.
- If the forwarding class is being moved to a local queue ID within the SAP ingress QoS policy and it is the first forwarding class to be mapped to the queue ID the system will attempt to instantiate the queue on each ingress SAP where the SAP ingress QoS policy is applied. If the queue cannot be created on any of the SAPs, the redirection removal from the current queue group ID will fail.

When a SAP egress QoS policy forwarding class is redirected to a queue group queue ID:

- If the queue group name does not exist as an egress queue group template, the forwarding class redirection will fail.
- If a redirection queue ID does not exist within the egress queue group template, the forwarding class redirection will fail.

- If the SAP egress QoS policy is currently applied to a non-Ethernet port or an Ethernet port where the specified egress queue group does not exist, the forwarding class redirection will fail.

When a SAP egress QoS policy forwarding class redirection is removed from a queue group queue ID:

- If the forwarding class is being moved to another queue group queue ID that does not exist within an egress queue group template, the redirection removal from the current queue group queue ID will fail.
- If the forwarding class is being moved to a local queue ID within the SAP egress QoS policy and the local queue ID does not exist, the redirection removal from the current queue group queue ID will fail.
- If the forwarding class is being moved to a local queue ID within the SAP egress QoS policy and it is the first forwarding class to be mapped to the queue ID the system will attempt to instantiate the queue on each egress SAP where the SAP egress QoS policy is applied. If the queue cannot be created on any of the SAPs, the redirection removal from the current queue group ID will fail.

If the above operation is successful then:

- The system decrements the association counter for the egress queue group template with the same name as the queue group previously specified in the forwarding class redirection.
- The system decrements the queue ID association counter within the queue group template for the queue ID previously specified in the forwarding class redirection.
- The system decrements the port queue group association counter for each egress port queue group where the SAP egress QoS policy is applied to a SAP.

When a SAP ingress QoS policy with a forwarding class redirection to a queue group queue ID is applied to a SAP:

- If the queue group specified in any forwarding class redirection does not exist as an ingress port queue group on the port associated with the SAP, the SAP ingress QoS policy application will fail.

If the operation above is successful, then:

- The system increments the port queue group association counter for each ingress port queue group referenced in a forwarding class redirection on the port associated with the SAP. The ingress port queue group association counter is incremented for each forwarding class redirected to the queue group within the added policy.

When a SAP ingress QoS policy with a forwarding class redirection to a queue group queue ID is removed from a SAP:

- If removing the SAP ingress QoS policy from the SAP results in the need to instantiate an ingress queue for the SAP that cannot be created, the SAP ingress QoS policy removal action will fail.

If the operation above is successful, then:

- The system decrements the port queue group association counter for each egress port queue group referenced in a forwarding class redirection within the removed SAP egress QoS policy. The egress port queue group association counter is decremented for each forwarding class redirected to the queue group within the removed policy.

SAP-Based Provisioning Model

When a redirection to a named forwarding plane queue group instance is applied to a SAP on ingress:

- If the queue group name does not exist as an ingress queue group template, the redirection will fail.
- If a queue group name does exist as an ingress queue group template, but the specified instance-id has not been instantiated on the same forwarding plane as used by the SAP, the redirection will fail.
- If a redirected policer ID in the SAP ingress QoS policy does not match a policer ID in the named ingress queue group template, the redirection will fail.
- If the SAP ingress QoS policy is currently applied to a non-Ethernet port or an Ethernet port where the specified ingress queue group instance does not exist on the forwarding plane, the redirection will fail.

If the operation above is successful, then:

- The system increments the association counter for the ingress queue group template with the same name as the queue group specified in the SAP redirection for each forwarding class redirected to the template.
- The system increments the policer ID association counter within the queue group template for each forwarding class redirected to a policer ID.
- The system increments the forwarding plane queue group instance association counter for each ingress queue group instance where a SAP ingress QoS policy specifying redirection is applied to a SAP.

When redirection to a named queue group is removed from an ingress SAP:

- If the forwarding class is being moved to another queue group policer ID that does not exist within the ingress FP queue group, the redirection removal from the current queue group policer ID will fail.

- If the forwarding class is being moved to a local policer ID within the SAP ingress QoS policy and the local policer ID does not exist, the redirection removal from the current queue group policer ID will fail.
- If the forwarding class is being moved to a local policer ID within the SAP ingress QoS policy and it is the first forwarding class to be mapped to the policer ID the system will attempt to instantiate the policer on each ingress SAP where the SAP ingress QoS policy is applied. If the policer cannot be created on any of the SAPs, the redirection removal from the current queue group policer ID will fail.

If the operation above is successful, then:

- The system decrements the association counter for the ingress queue group template with the same name as the queue group previously specified in the forwarding class redirection.
- The system decrements the policer ID association counter within the queue group template for the policer ID previously specified in the forwarding class redirection.

The system decrements the forwarding plane queue group template association counter for each ingress queue group where redirection is applied to the ingress SAP.

For the SAP-based provisioning model, the rules for redirecting a forwarding class queue to an egress port queue group are similar to those on ingress.

- If an egress QoS policy containing one or more redirections is applied to a SAP, but either no queue group instance is specified at association time, or a named queue group instance is specified and either the queue group name or the instance identifier does not correspond to a queue group that has been created on the egress port, then the association will be rejected.
- If all of the redirections in an egress QoS policy are to queue ids that do not exist in the named queue group instance, then the association will be rejected.
- Note that if a policer local to a SAP feeds into a SAP based queue group queue instance, and the queue ID to use is not explicitly specified in the egress QoS policy (through the command policer policer-id port-redirect-group-queue) and is instead inferred from the forwarding class of the policer, but that forwarding class does not exist in the queue group template, then no error is generated. Instead, the queue with the lowest queue ID is used in the queue group instance. If at a later time, a user attempts to add a queue with a given queue-id to a policer redirect for a given forwarding class in the egress QoS template, then the system will check that the corresponding queue-id exists in any queue group instances associated with any SAPs using the QoS policy.

Access Queue Group Statistics

Port Queue Groups

When a forwarding class is redirected to a ingress or egress port queue group queue, the packets sent to the queue are statistically tracked by a set of counters associated with the queue group queue and not with any of the counters associated with the SAP.

This means that it is not possible to perform accounting within a queue group based on the source SAPs feeding packets to the queue. The statistics associated with the SAP will not reflect packets redirected to a port queue group queue.

The set of statistics per queue are eligible for collection in a similar manner as SAP queues. The collect-stats command enables or disables statistics collection in to a billing file based on the accounting policy applied to the queue group.

Forwarding Plane Queue Groups

When a forwarding class is redirected to a forwarding plane queue group queue or policer, the packets sent to the queue or policer are statistically tracked by a set of counters associated with the queue group queue/policer and not with any of the counters associated with the SAP.

This means that it is not possible to perform accounting within a queue group based on the source SAPs feeding packets to the queue. That is, the statistics associated with the SAP will not include packets redirected to a queue group queue.

Note that if the user enables the **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*} option under the ingress queue-group policer, the byte counters of that policer will reflect the adjusted packet size.

The set of statistics per queue are eligible for collection in a similar manner to SAP queues. The **collect-stats** command enables or disables statistics collection in to a billing file based on the accounting policy applied to the queue group.

Network IP Interface Forwarding Class-Based Redirection

Forwarding class redirection for a network IP interface is defined in a four step process.

1. Create an ingress or egress queue group template with the appropriate queues or policers.
2. Apply an instance of an ingress queue-group template created in step 1 (containing only policers) to the FP ingress network configuration context of card X. In addition, or alternatively, apply an instance of an egress queue-group template created in step 1 to the network egress configuration context of port Y.
3. Configure the network QoS policy used on the IP interface to redirect ingress traffic to a policer ID (defined in the ingress queue-group template created in step 1) on the basis of forwarding-class and forwarding-type (unicast vs. multicast). In addition, or alternatively, configure the network QoS policy to redirect egress traffic to a queue ID and/or a policer ID based on forwarding-class.
4. Apply the network QoS policy to the network IP interface and at the same time specify the ingress and/or egress queue-group instances associated with the interface.

Egress Network Forwarding Class Redirection Association Rules

The association rules work differently for network egress IP interfaces than they do for access SAPs. Since the network QoS policy does not directly reference the queue group names, the system is unable to check for queue group template existence or queue ID existence when the forwarding class queue redirection is defined. Configuration verification can only be checked at the time the network QoS policy is applied to a network IP interface.

The system keeps an association counter for each queue group template and an association counter for each queue ID within the template. The system also keeps an association counter for each queue group created on a port.

When a network QoS policy is applied to an IP interface with the queue group parameter specified:

- If the queue group name does not exist as an egress queue group template, the QoS policy application will fail.
- If a redirection queue ID within the policy does not exist within the egress queue group template, the QoS policy application will fail.
- If the IP interface is bound to a port (or LAG) and the specified queue group name does not exist on the port, the QoS policy application will fail.

If the operation above is successful, then:

- The system increments the association counter for the queue group template with the same name as the queue group specified when the QoS policy is applied.
- The system increments the queue ID association counter within the queue group template for each forwarding class redirected to the queue ID.
- If the IP interface is currently bound to a port (or LAG), the association counter for the queue group on the port is incremented.

When the queue group parameter is removed from an IP interface:

- The system decrements the association counter for the queue group template with the same queue group name that was removed from the IP interface.
- The system decrements the queue ID association counter within the queue group template for each forwarding class that had previously been redirected to the queue ID.
- If the IP interface is currently bound to a port (or LAG), the association counter for the removed queue group on the port is decremented.

When a network QoS policy egress forwarding class redirection to a queue ID is removed or added:

- If a redirection is being added to a forwarding class and the queue ID does not exist on the queue groups for IP interfaces where the QoS policy is applied, the redirection will fail.

If the operation above is successful, then:

- The system finds all IP interfaces where the policy is applied.
- Finds all affected queue group templates based on the queue group associated with the QoS policy on each interface.
- If removing, the queue ID association counter is decremented within each queue group template based on the queue ID removed from the policy.
- If adding, the queue ID association counter is incremented within each queue group template based on the queue ID added to the policy.

When an IP interface associated with a queue group is bound to a port:

- If the specified egress queue group does not exist on the port, the port binding will fail.

If the operation above is successful, then:

- The system increments the association counter for the queue group on the port.

When an IP interface associated with a queue group is unbound from a port:

- The system decrements the association counter for the queue group on the unbound port

Egress Network IP Interface Statistics

The statistics for network interfaces work differently than statistics on SAPs. Counter sets are created for each egress IP interface and not per egress queue. When a forwarding class for an egress IP interface is redirected from the default egress port queue to a queue group queue, the system continues to use the same counter set.

Separate Ingress IPv4 and IPv6 Statistics

This feature adds support for separate ingress IPv4 and IPv6 statistics on IP interfaces. IES and VPRN interfaces, and subscriber group interfaces on IES and VPRN, as well as for uRPF. In previous release, the ingress statistics for IPv4 and IPv6 traffic was combined into a single set of packet and bytes counters. The existing counters will now only count IPv4 traffic, while new separate counters are available to IPv6 traffic.

The feature introduces a new CLI command to explicitly enable ingress statistics on IP interfaces, changing the default to disabled.

Ingress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection

Feature Configuration

The user applies a network QoS policy to the ingress context of a spoke-SDP[1] to redirect the mapping of a Forwarding Class (FC) to a policer defined in a queue-group template which is instantiated on the ingress Forwarding Plane (FP) where the PW packets are received.

config>service>vprn>interface>spoke-sdp>ingress>qos *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*

Let us refer to a queue-group containing policers as a *policer queue-group*. The user must instantiate this queue-group by applying the following command:

config>card>fp>ingress>network>queue-group *queue-group-name* **instance** *instance-id*

The policers are instantiated at ingress FP, one instance per destination tap, and are used to service packets of this spoke-SDP which are received on any port on the FP to support a network IP interface on LAG and on any network IP interface to support ECMP on the network IP interface and LSP reroutes to a different network IP interface on the same FP.

In the ingress context of the network QoS policy, the user defines the mapping of a FC to a policer-id and instructs the code to redirect the mapping to the policer of the same ID in some queue-group:

config>qos>network>ingress>fc>fp-redirect-group policer *policer-id*
config>qos>network>ingress>fc>fp-redirect-group multicast-policer *policer-id*

The user can redirect the Unicast and multicast packets of a FC to two different policers to allow for different policing rates for these packet types. However, the queue-group is explicitly named only at the time the network QoS policy is applied to the spoke-SDP as shown above with the example of the VPRN service.

When the FC of a PW is redirected to use a policer in the named queue-group, the policer feeds the existing per-FP ingress shared queues referred to as *policer-output-queues*. These queues are shared by both access and network policers configured on the same ingress FP. The shared queue parameters are configurable using the following command:

configure>qos>shared-queue policer-output-queues

1. This feature applies to both spoke-SDP and mesh-SDP. Spoke-SDP is used throughout for ease of reading.

The CLI configuration in this section uses a spoke-SPD defined in the context of a VPRN interface. However the PW shaping feature is supported with all PW based services including the PW template.

Provisioning Model

Operationally, the provisioning model in the case of the ingress PW shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC which needs to be redirected and optionally for each traffic type (Unicast or multicast).
2. Apply the queue-group template to the network ingress context of all IOM3/IMM FPs where there exists a network IP interface which the PW packets can be received on. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step which means the same network QoS policy can redirect different PWs to different queue-group templates.
    a. Apply this network QoS policy to the ingress context of a spoke-SDP inside a service or to the ingress context of a PW template and specify the redirect queue-group name.
    b. One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress PW shaping feature:

1. Only a queue-group containing policers, can be instantiated in the network ingress context of an IOM3/IMM FP. If the queue-group template contains policers and queues, the queues are not instantiated.
2. If the queue-group contains queues only, the instantiation in the data path is failed.
3. One or more instances of the same policer queue-group name and/or a different policer queue-group name can be created on network ingress context of an IOM3/IMM FP.
4. The queue-group-name must be unique within all network ingress and access ingress queue groups in the system.
5. The instantiated FP policer queue-group can be used by PW packets received on a network IP interface configured on both Ethernet ports and POS ports of that IOM3/IMM.

6. When a PW FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the PW packet feeds directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
7. When a PW FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the PW packet feeds directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
8. When a PW FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:
    a. When a PW packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and is then feed the per-FP ingress shared queues referred to as *policer-output-queues*.
    b. When a PW packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the PW packets are fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
9. If a network QoS policy is applied to the ingress context of a PW, any PW FC, which is not explicitly redirected in the network QoS policy, has the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
    a. This behavior is the same regardless if the ingress network IP interface from which the PW packet is received is redirected or not to a policer queue-group.
10. If no network QoS policy is applied to the ingress context of the PW, then all packets of the PW feed:
    a. the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP. This is the default behavior.
    b. a queue-group policer followed by the per-FP ingress shared queues referred to as *policer-output-queues* if the ingress context of the network IP interface from which the packet is received is redirected to a queue-group. The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from a R-VPLS spoke-SDP, which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP is used.

Ingress Packet Classification

When a PW is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the PW. This is true regardless if an instance of the named policer queue-group exists on the ingress FP the PW packet is received on. The user can apply a QoS filter matching the dot1.p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload's IP header if the user enabled the **ler-use-dscp** option and the PW terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the PW is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface the PW packet is received on.

Egress PW Shaping using Spoke-SDP Forwarding Class-Based Redirection

Feature Configuration

The user applies a network QoS policy to the egress context of a spoke-sdp to redirect the mapping of a Forwarding Class (FC) to a policer and/or a queue part of a queue-group instance created in the egress of a network port.

config>service>vprn>interface>spoke-sdp>egress>qos *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

The queue-group queues or policers are instantiated at egress port, one instance per network port and per link of LAG network port and are used to service packets of this spoke-SDP, which are forwarded over any network IP interface on this port.

config>port>ethernet>network>egress>queue-group *queue-group-name* **instance** *instance-id*

In the egress context of the network QoS policy, the user defines the mapping of a FC to a policer-id or a queue-id and instructs the code to redirect the mapping to the queue or policer of the same ID in some queue-group. However, the queue-group is explicitly named only at the time the network QoS policy is applied to the spoke-SDP as shown above with the example of the VPRN service. The command is as follows:

config>qos>network>egress>fc>port-redirect-group {**queue** *queue-id* | **policer** *policer-id* [**queue** *queue-id]*}

There are three possible outcomes when executing this command.

- The user can redirect a FC to use a queue in a queue-group and in which case there are no policers used.
  config>qos>network>egress>fc>port-redirect-group queue *queue-id*
- The user can redirect a FC to use a policer-id in a queue-group without specifying a queue-id and in which case the policer is feeding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.
  config>qos>network>egress>fc>port-redirect-group policer *policer-id*
- The user can redirect a FC to use a policer feeding a queue both of which are defined in the named queue-group.
  config>qos>network>egress>fc>port-redirect-group policer *policer-id*
  queue *queue-id*

The CLI configuration in this section uses a spoke-sdp defined in the context of a VPRN interface. However the PW shaping feature is supported with all PW based services and PW template.

Provisioning Model

This provisioning model allows the mapping of one ore more PWs to the same instance of queues, or policers and queue, which are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only or policers and queues for each FC which needs to be redirected.
2. Apply the queue-group template to the network egress context of all IOM3/IMM ports where there exists a network IP interface which the PW packets can be forwarded on. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step which means the same network QoS policy can redirect different PWs to different queue-group templates.
    a. Apply this network QoS policy to the egress context of a spoke-sdp inside a service or to the egress context of a PW template and specify the redirect queue-group name.
    b. One or more spoke-sdp's can have their FCs redirected to use queues only or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model:

1. Queue-groups containing queues only or policers and queues can be instantiated in the network egress context of an Ethernet port on IOM3/IMM.
2. When a port is a LAG, one instance of the queue-group is instantiated on each member link.
3. One or more instances of the same queue-group name and/or a different queue-group name can be created in the network egress context of an Ethernet port.
4. The queue-group-name must be unique within all network egress and access egress queue groups in the system.
5. A user attempt to instantiate the queue-group on the network egress context of a POS port or a TDM port will fail.
6. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the PW packet is fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on. This queue can be a queue-group

queue or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a PW packet.

7. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the PW packet is fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on.

8. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports which have network IP interfaces. The handling of this is dealt with in the data path as follows:

    a. When a PW packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and is fed to the queue-group queue. If only a policer is specified in the redirection command, then the packet is processed by the queue-group policer and is then fed into the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. If only a queue is specified in the redirection command, the packet is fed to the queue-group queue.

    b. When a PW packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the PW packet is fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

    c. If a network QoS policy is applied to the egress context of a PW, any PW FC, which is not explicitly redirected in the network, QoS policy has the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

Egress Marking of PW Packet Header

When the queue-group name the PW is redirected to exists and the redirection succeeds, the marking of the packet's DEI/dot1.p/DSCP and the tunnel's DEI/dot1.p/DSCP/EXP is performed according to the relevant mappings of the {FC, profile} in the egress context of the network QoS policy applied to the PW. This is true if an instance of the queue-group exists on the egress port the PW packet is forwarded to. If the packet's profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet's DEI/dot1.p and the tunnel's DEI/dot1.p/EXP but the DSCP is not modified by the policer's operation.

When the redirection command succeeds but there is no instance of the queue-group on the egress port, or when the redirection command fails due to an inexistent queue-group name, the marking of the packet's DEI/dot1.p/DSCP and the tunnel's DEI/dot1.p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface the PW packet is forwarded to.

Egress Packet Re-Classification Based on IPv4/IPv6 Criteria

The user enables IP precedence or DSCP based egress re-classification by applying the following command in the context of the network QoS policy applied to the egress context of a spoke-SDP.

config>qos>network>egress>prec *ip-prec-value* [**fc** *fc-name*] [**profile** {**in** | **out**}]
config>qos>network>egress>dscp *dscp-name* [**fc** *fc-name*] [**profile** {**in** | **out**}]

The IP precedence bits used to match against DSCP reclassification rules come from the Type of Service (ToS) field within the IPv4 header or the Traffic Class field from the IPv6 header.

The IP DSCP bits used to match against DSCP reclassification rules come from the Type of Service (ToS) field within the IPv4 header or the Traffic Class field from the IPv6 header.

If the packet does not have an IP header, DSCP or IP-precedence based matching is not performed.

Note that the IP precedence and DSCP based re-classification are only supported on a PW used in an IES or VPRN spoke-interface. The CLI blocks the application of a network QoS policy with the egress re-classification commands to a network IP interface or to a spoke-SDP part of L2 service. Conversely, the CLI does not allow the user to add the egress re-classification commands to a network QoS policy if it is being used by a network IP interface or a L2 spoke-SDP.

In addition, the egress re-classification commands only take effect if the redirection of the spoke-SDP to use an egress port queue-group succeeds; for example, the following CLI commands succeed:

config>service>vprn>interface>spoke-sdp>egress>qos *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

config>service>ies>interface>spoke-sdp>egress>qos *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

Reclassification will however occur regardless of whether the queue group instance exists or not on a given egress network port. When the redirection command fails in CLI, the PW uses the network QoS policy assigned to the network IP interface. Since the network QoS policy applied to a network IP interface does not support re-classification, the PW packets do not undergo re-classification.

Ingress and Egress PW Statistics

The PW forwarded packet and octet statistics (SDP binding statistics) are currently supported for both ingress and egress and are available via show command, monitor command, and accounting file. These statistics consist of the ingress-forwarded and ingress-dropped packet and octet counters, as well as the egress-forwarded packet and octet counters. However, they do not include discards in the ingress network queues. The latter are counted in the stats of the queues defined in the network-queue policy applied to the ingress of the MDA/FP.

Note the ingress and egress SDP binding stats do not count the label stack of the PW packet but count the PW Control Word (CW) if included in the packet.

With the introduction of the PW shaping feature—the ingress or egress queue-group policer—a PW FC is redirected to also provide packet and octet forwarded and dropped-statistics by means of the show command, monitor command, and accounting file of the ingress or egress queue-group instance.

Similar to the SDP binding stats, the ingress policer stats for a spoke-SDP does not count the label stack. When the spoke-SDP is part of a L2-service, they will count the L2-encapsulation, minus CRC and VLAN tag if popped out, and they also count the PW CW, if included in the packet. When the spoke-SDP is part of a L3-service, the policer stats only count the IP payload and do not count the PW CW. Unlike the ingress SDP binding stats, if the user enables the **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*} option under the queue-group policer, then the policer stats reflect the adjusted packet size in both L2 and L3-spoke-SDPs.

The egress queue-group policer and/or queue counts the full label stack of the PW packet including the CW. If the user enables the **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*} option under the queue-group policer and queue-group queue, then the policer and queue stats reflect the adjusted packet size.

The SDP binding and queue-group statistics does however remain separate as one or more PWs can have FCs redirected to the same policer ID in the queue-group instance.

Queue Group Behavior on LAG

Queue Group Queue Instantiation Per Link

When a port queue group is created on a Link Aggregation Group (LAG) context, it is individually instantiated on each link in the LAG.

Per Link Queue Group Queue Parameters

The queue parameters for a queue within the queue group are used for each port queue and are not divided or split between the port queues representing the queue group queue. For instance, when a queue rate of 100Mbps is defined on a queue group queue, each instance of the queue group (on each LAG port) will have a rate of 100Mbps.

Adding a Queue Group to an Existing LAG

A queue group must be created on the primary (lowest port ID) port of the LAG. If an attempt is made to create a queue group on a port other than the primary, the attempt will fail. When the group is define on the primary port, the system will attempt to create the queue group on each port of the LAG. If sufficient resources are not available on each port, the attempt to create the queue group will fail.

Any queue group queue overrides defined on the primary port will be automatically replicated on all other ports within the LAG.

Removing a Queue Group from a LAG

A queue group must be removed from the primary port of the LAG. The queue group will be deleted by the system from each of the port members of the LAG.

Adding a Port to a LAG

When adding a port to a LAG group, the port must have the same queue groups defined as the existing ports on the LAG before it will be allowed as a member. This includes all queue group override parameters.

Basic Configurations

- Configuring an Ingress Queue Group Template on page 413
- Configuring Egress Queue Group Template on page 414
- Applying Ingress Queue Group to SAP Ingress Policy on page 415
- Applying Egress Queue Group to SAP Egress Policy on page 416
- Configuring Queue Group on Ethernet Access Ingress Port on page 419
- Configuring Overrides on page 421
- Configuring Queue Group on Ethernet Access Egress Port on page 422
- Configuring Queue Group for Network Egress Traffic on Port on page 423
- Configuring Queue Group for Network Ingress Traffic on Forwarding Plane on page 424
- Using Queue Groups to Police Ingress/Egress Traffic on Network Interface on page 425
- Configuring Ingress/Egress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection on page 427
- Specifying QoS Policies on Service SAPs on page 430

Configuring an Ingress Queue Group Template

The following displays an ingress queue group template configuration example:

```
*A:Dut-T>cfg>qos>qgrps# info
----------------------------------------------
            ingress
                queue-group "QG_ingress_1" create
                    queue 1 best-effort create
                        mbs 100
                    exit
                    queue 2 best-effort create
                        mbs 100
                    exit
                    queue 3 best-effort create
                        mbs 100
                    exit
                    queue 4 best-effort create
                        mbs 100
                    exit
                exit
            exit
----------------------------------------------
*A:Dut-T>cfg>qos>qgrps#
```

NOTE: To fully use the queue group feature to save queues, you must explicitly map all forwarding classes to queue group queues. This rule is applicable to SAP ingress, SAP egress and network QoS policies.

Configuring Egress Queue Group Template

The following displays an egress queue group template configuration example:

```
*A:Dut-T>cfg>qos>qgrps# info
----------------------------------------------
...
            egress
                queue-group "QG_egress_1" create
                    description "Egress queue group"
                    queue 1 best-effort create
                        mbs 100
                    exit
                    queue 2 best-effort create
                        mbs 100
                    exit
                    queue 3 best-effort create
                        mbs 100
                    exit
                    queue 4 best-effort create
                        mbs 100
                    exit
                exit
            exit
----------------------------------------------
*A:Dut-T>cfg>qos>qgrps#
```

Applying Ingress Queue Group to SAP Ingress Policy

The following display a SAP ingress policy configuration with **group** *queue-group-name* specified:

```
*A:Dut-T>config>qos>sap-ingress# info
----------------------------------------------
            queue 1 create
            exit
            queue 11 multipoint create
            exit
            fc "af" create
                queue 2 group "QG_ingress_1"
            exit
            fc "be" create
                queue 1 group "QG_ingress_1"
            exit
            fc "ef" create
                queue 3 group "QG_ingress_1"
            exit
            fc "nc" create
                queue 4 group "QG_ingress_1"
            exit
            dot1p 0 fc "be"
            dot1p 2 fc "af"
            dot1p 4 fc "ef"
            dot1p 6 fc "nc"
----------------------------------------------
*A:Dut-T>config>qos>sap-ingress#
```

Applying Egress Queue Group to SAP Egress Policy

The following display a SAP egress policy configuration with **group** *queue-group-name* specified:

```
A:Dut-T>config>qos>sap-egress# info
----------------------------------------------
            queue 1 create
            exit
            fc af create
                queue 2 group "QG_egress_1"
            exit
            fc be create
                queue 1 group "QG_egress_1"
            exit
            fc ef create
                queue 3 group "QG_egress_1"
            exit
            fc nc create
                queue 4 group "QG_egress_1"
            exit
----------------------------------------------
A:Dut-T>config>qos>sap-egress#
```

SAP-based Egress Queue Re-direction

The following displays a SAP egress policy configuration with port-redirect-group-queue construct (shown for both regular and HS-MDA egress queues) and the actual queue-group-name is determined by the SAP egress QoS configuration:

```
*A:Dut-A# configure qos sap-egress 3
*A:Dut-A>config>qos>sap-egress# info
----------------------------------------------
            queue 1 create
            exit
            queue 2 create
            exit
            policer 8 create
                rate 50000
            exit
            fc af create
                queue 3 port-redirect-group-queue
                hsmda
                    queue 3 port-redirect-group-queue
                exit
            exit
            fc be create
                queue 3 port-redirect-group-queue
                hsmda
                    queue 3 port-redirect-group-queue
                exit
            exit
            fc ef create
                policer 8 port-redirect-group-queue
                hsmda
                    queue 3 port-redirect-group-queue
                exit
            exit
            fc h1 create
                queue 3 port-redirect-group-queue
                hsmda
                    queue 3 port-redirect-group-queue
                exit
            exit
            fc h2 create
                queue 3 port-redirect-group-queue
                hsmda
                    queue 3 port-redirect-group-queue
                exit
            exit
            fc l1 create
                queue 3 port-redirect-group-queue
                hsmda
                    queue 3 port-redirect-group-queue
                exit
            exit
            fc l2 create
                queue 3 port-redirect-group-queue
                hsmda
```

```
                    queue 3 port-redirect-group-queue
                exit
            exit
            fc nc create
                queue 3 port-redirect-group-queue
                hsmda
                    queue 3 port-redirect-group-queue
                exit
            exit
-----------------------------------------------

This is to be in-conjunction with:

*A:Dut-A# configure service vpls 1
*A:Dut-A>config>service>vpls# info
-----------------------------------------------
            stp
                shutdown
            exit
            sap 9/1/2:1 create
                egress
                    qos 3 port-redirect-group qg1 instance 101
                exit
            exit
```

Configuring Queue Group on Ethernet Access Ingress Port

The provisioning steps involved in using a queue-group queue on an ingress port are:

- Queue Group Template Creation
  - → Create the queue group template in the ingress context
  - → Create the queue within the queue group template
- Queue Group Creation
  - → Identify the ingress port (or ports) for which the queue group will be needed (for LAG use the primary port member)
  - → Create a queue group with the same name as the template on the port or ports
- Map a Forwarding Class to the queue-id within the queue group
  - → Map forwarding classes to queue-group queues.
  - → Identify or create the SAP ingress QoS policy that will be used on the ingress SAP where queue redirection is desired
  - → Map the desired forwarding classes to the queue group name and the specific queue ID within the group
- Apply the SAP ingress QoS policy
  - → Identify or create the ingress SAP requiring forwarding class redirection to the queue group
  - → Assign the QoS policy to the SAP

The following displays an Ethernet access ingress port queue-group configuration example:

```
*A:Dut-T>config>port# /configure port 9/2/1
*A:Dut-T>config>port# info
----------------------------------------------
        ethernet
            mode access
            access
                ingress
                    queue-group "QG_ingress_1" create
                    exit
                exit
                egress
                    queue-group "QG_egress_1" create
                    exit
                exit
            exit
        exit
        no shutdown
----------------------------------------------
*A:Dut-T>config>port#
```

```
*A:Dut-T>config>port# /configure port 9/2/2
*A:Dut-T>config>port# info
----------------------------------------------
        ethernet
            mode access
            access
                ingress
                    queue-group "QG_ingress_1" create
                    exit
                exit
                egress
                    queue-group "QG_egress_1" create
                    exit
                exit
            exit
        exit
        no shutdown
----------------------------------------------
*A:Dut-T>config>port#
```

Configuring Overrides

The following output display a port queue group queue override example.

```
*A:Dut-T>config>port>ethernet>access# /configure port 9/2/1
*A:Dut-T>config>port# info
----------------------------------------------
        ethernet
            mode access
            access
                ingress
                    queue-group "QG_ingress_1" create
                        queue-overrides
                            queue 2 create
                                rate 800000 cir 20000
                            exit
                        exit
                    exit
                exit
                egress
                    queue-group "QG_egress_1" create
                    exit
                exit
            exit
        exit
        no shutdown
----------------------------------------------
*A:Dut-T>config>port# /configure port 9/2/2
*A:Dut-T>config>port# info
----------------------------------------------
        ethernet
            mode access
            access
                ingress
                    queue-group "QG_ingress_1" create
                    exit
                exit
                egress
                    queue-group "QG_egress_1" create
                        queue-overrides
                            queue 3 create
                                rate 1500000 cir 2000
                            exit
                        exit
                    exit
                exit
            exit
        exit
        no shutdown
----------------------------------------------
*A:Dut-T>config>port#
```

Configuring Queue Group on Ethernet Access Egress Port

The provisioning steps involved in using a queue-group queue on an egress access port are:

- Queue Group Template Creation
  - → Create the queue group template in the egress context
  - → Create the queue within the queue group template
- Queue Group Creation
  - → Identify which egress port (or ports) on which the queue group will be needed (for LAG use the primary port member)
  - → Create a queue group instance with the same name as the template on the port or ports
- From this point, there are two methods for regular ethernet based SAPs to have port access egress re-direction. a). Policy based re-direction and, b). SAP based re-direction. For Policy based redirection:
- Map a Forwarding Class to the queue-id within the queue group
  - → Identify or create the SAP egress QoS policy that will be used on the egress SAP where policy-based queue re-direction is desired
  - → Map the desired forwarding classes to the queue group name and the specific queue ID within the group with the “group” keyword
- Apply the SAP egress QoS policy
  - → Identify or create the egress SAP requiring forwarding class redirection to the queue group
  - → Assign the QoS policy to the SAP
- For SAP based redirection:
- Map a Forwarding Class to the queue-id within the queue group
  - → Identify or create the SAP egress QoS policy that will be used on the egress SAP where SAP-based queue re-direction redirection is desired
  - → Map the desired forwarding classes to the queue group specific queue-id, and the keyword "port-redirect-group-queue". The actual queue-group template name is determined by the sap instance's configuration which associated the sap-egress qos policy in conjunction with the port-redirect-group's instance.
- Apply the SAP egress QoS policy and the queue-group template's instance under the SAP.
  - → Identify or create the egress SAP requiring forwarding class redirection to the queue group
  - → Assign the QoS policy and the egress queue-group template's instance to the SAP.

Configuring Queue Group for Network Egress Traffic on Port

The provisioning steps involved in using a queue-group queue on an egress network port are:

- Queue Group Template Creation:
  - → Create the egress queue group template.
  - → Create the queues and/or policers within the queue group template.
- Queue Group Creation:
  - → Identify the egress port (or ports) on which the queue group will be needed (for LAG use the primary port member).
  - → Create a queue group with the same name as the template on the port or ports. The instance ID is optional.
- Map a Forwarding Class to the queue-id within the queue group:
  - → Identify or create the network QoS policy that will be used on the egress IP interface where queue redirection is desired.
  - → Map the desired egress forwarding classes within the network QoS policy to the specific queue IDs and/or policer IDs within the group (the group name will be supplied when the QoS policy is applied to the IP interface).
- Apply the network QoS policy:
  - → Identify or create the IP interface requiring forwarding class redirection to the queue group.
  - → Assign the QoS policy to the IP interface and specify the queue group name (and optionally instance ID) for redirection of egress traffic.

Once a queue within a template is mapped by a forwarding class on any object, the queue may be edited, but not deleted.

Configuring Queue Group for Network Ingress Traffic on Forwarding Plane

The provisioning steps involved in using a queue-group for ingress traffic on a network interface are:

- Queue Group Template Creation:
  - → Create the ingress queue group template.
  - → Create the policers within the queue group template.
- Queue Group Creation:
  - → Identify the ingress forwarding plane on which the queue group will be needed.
  - → Create a queue group with the same name as the template in the FP ingress network configuration context. An instance ID is mandatory.
- Map a Forwarding Class to the policer-id within the queue group:
  - → Identify or create the network QoS policy that will be used on the ingress IP interface where queue redirection is desired.
  - → Map the desired ingress forwarding classes within the network QoS policy to the specific policer IDs within the group (the group name will be supplied when the QoS policy is applied to the IP interface).
- Apply the network QoS policy:
  - → Identify or create the IP interface requiring forwarding class redirection to the queue group.
  - → Assign the QoS policy to the IP interface and specify the queue group name and instance ID for redirection of ingress traffic.

Using Queue Groups to Police Ingress/Egress Traffic on Network Interface

```
config
    qos
        queue-group-templates
            ingress
                queue-group "Ingress_QG_1" create
                    policer 2 create
                        rate 9000
                    exit
                exit
            exit
            egress
                queue-group "Egress_QG_1" create
                    queue 1 best-effort create
                    exit
                    policer 2 create
                        rate 9000
                    exit
                exit
            exit
        exit

        network 2 create
            ingress
                fc be
                    fp-redirect-group policer 2
                exit
            exit
            egress
                fc be
                    port-redirect-group policer 2
                exit
            exit
        exit


    card 1
        card-type xcm-x20
            mda 1              mda-type cx20-10g-sfp no shutdown
            exit
        fp 1
            ingress
                network
                    queue-group "Ingress_QG_1" instance 550 create
                    exit
                exit
            exit
        exit
        no shutdown

    port 1/1/3
        ethernet
                mtu 1514
                network
                    egress
                        queue-group "Egress_QG_1" instance 550 create
```

```
                        exit
                    exit
                exit
            exit
        no shutdown
exit

router
    interface “to-D”
    address 10.10.11.3/24
    port 1/1/3
    qos 2 egress-port-redirect-group "Egress_QG_1" egress-instance
    550 ingress-fp-redirect-group "Ingress_QG_1" ingress-instance
    550
    no shutdown
```

Configuring Ingress/Egress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection

```
configure
#--------------------------------------------------
echo "QoS Policy Configuration"
#--------------------------------------------------
    qos
        queue-group-templates
            ingress
                queue-group "QGIng1" create
                    policer 1 create
                    exit
                    policer 2 create
                    exit
                    policer 3 create
                    exit
                    policer 4 create
                    exit
                exit
            exit
            egress
                queue-group "QGEgr1" create
                    queue 1 best-effort create
                    exit
                    policer 1 create
                    exit
                    policer 2 create
                    exit
                    policer 3 create
                    exit
                    policer 4 create
                    exit
                exit
            exit
        exit
    exit
        network 10 create
            ingress
                lsp-exp 0 fc be profile out
                lsp-exp 1 fc be profile out
                lsp-exp 2 fc be profile out
                lsp-exp 3 fc be profile out
                lsp-exp 4 fc be profile out
                lsp-exp 5 fc be profile out
                lsp-exp 6 fc be profile out
                lsp-exp 7 fc be profile out
                fc af
                    fp-redirect-group policer 4
                exit
                fc be
                    fp-redirect-group policer 1
                exit
                fc l1
                    fp-redirect-group policer 2
                exit
                fc l2
```

```
                    fp-redirect-group policer 3
                exit
            exit
            egress
                fc af
                    port-redirect-group policer 4
                exit
                fc be
                    port-redirect-group policer 1
                exit
                fc l1
                    port-redirect-group policer 2
                exit
                fc l2
                    port-redirect-group policer 3
                exit
            exit
        exit
    exit
#--------------------------------------------------
echo "Card Configuration"
#--------------------------------------------------
    card 3
        fp 1
            ingress
                network
                    queue-group "QGIng1" instance 1 create
                    exit
                    queue-group "QGIng1" instance 2 create
                    exit
                exit
            exit
        exit
    exit
#--------------------------------------------------
echo "Port Configuration"
#--------------------------------------------------
    port 3/2/1
        ethernet
            encap-type dot1q
            network
                egress
                    queue-group "QGEgr1" instance 1 create
                    exit
                    queue-group "QGEgr1" instance 2 create
                    exit
                exit
            exit
        exit
        no shutdown


*A:Dut-T>config>service#
        customer 1 create
            description "Default customer"
        exit
        sdp 1 mpls create
            description "Default sdp description"
            far-end 2.2.2.2
```

```
            ldp
            path-mtu 9000
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        vpls 1 customer 1 vpn 1 create
            description "Default tls description for service id 1"
            service-mtu 9000
            stp
                shutdown
            exit
            service-name "XYZ Vpls 1"
            sap 9/2/1:1.* create
                description "Default sap description for service id 1"
                static-mac 00:00:1e:00:01:02 create
                ingress
                    qos 10
                exit
            exit
            spoke-sdp 1:101 vc-type vlan create
                description "Description for Sdp Bind 1 for Svc ID 1"
                ingress
                    qos 10 fp-redirect-group "QGIng1" instance 1
                exit
                egress
                    qos 10 port-redirect-group "QGEgr1" instance 1
                exit
                static-mac 00:00:28:00:01:02 create
                no shutdown
            exit
            no shutdown
        exit


    router
        interface "ip-12.1.1.1"
            address 12.1.1.1/24
            port 3/2/1:1
        exit
        interface "system"
            address 1.1.1.1/32
        exit
#--------------------------------------------
```

Specifying QoS Policies on Service SAPs

The following output displays a VPLS service configuration example.

```
*A:Dut-T>config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            sap 9/2/1 create
                ingress
                    qos 10
                exit
                egress
                    qos 10
                exit
            exit
            sap 9/2/2 create
                ingress
                    qos 10
                exit
                egress
                    qos 10
                exit
            exit
            no shutdown
----------------------------------------------
*A:Dut-T>config>service>vpls#
```

Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
no description

Context cfg>qos>qgrps>egr>qgrp
cfg>qos>qgrps>ing>qgrp
cfg>qos>qgrps>ing>qgrp>policer

Description This command creates a text description stored in the configuration file for a configuration context. The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default none

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Queue Group Commands

queue-group-templates

Syntax **queue-group-templates**

Context config>qos

Description This command enables the context to define ingress and egress queue group templates.

Default none

egress

Syntax **egress**

Context cfg>qos>qgrps

Description This command enables the context to configure QoS egress queue groups. Egress queue group templates can be applied to egress Ethernet ports to create an egress queue group.

Default none

queue-group

Syntax **queue-group** *queue-group-name* [**create**]
no queue-group *queue-group-name*

Context cfg>qos>qgrps>egr
cfg>qos>qgrps>ingr

Description This command creates a queue group template. The system does not maintain default queue groups or queue group templates. Each queue group template used in the system must be explicitly created.

The *queue-group-name* parameter is required when executing the queue-group command and identifies the name of the template to be either created or edited. Each ingress queue group template must be uniquely named within the system. Multiple ingress queue group templates may not share the same name. An ingress and egress queue group template may share the same name.

The **no** form of the command removes the specified queue group template from the system. If the queue group template is currently in use by an ingress port, the command will fail. If group-name does not exist, the command has no effect and does not return an error.

Default none

Parameters *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.

create — Keyword used to create the queue group instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

policer

Syntax **policer** *policer-id* [**create**]
no policer *policer-id*

Context config>qos>queue-group-templates>ingress>queue-group
config>qos>queue-group-templates>egress>queue-group

Description This command is used in ingress and egress queue-group templates to create, modify, or delete a policer.

Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The ingress queue-group template may have up to 32 policers (numbered 1 through 32) and may be defined, while the egress queue-group template supports a maximum of 8 (numbered 1 through 8). While a policer may be defined in a queue-group template, it is not actually created until the queue-group template is instantiated on the ingress context of a forwarding plane or on the egress context of a port.

Once a policer is created, the policer's metering rate and profiling rates may be defined, as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues that have dedicated counters, policers allow various stat-mode settings which define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

Once a policer is created, it cannot be deleted from the queue-group template unless any forwarding classes that are redirected to the policer are first removed.

The **no** version of this command deletes the policer.

Parameters *policer-id* — The *policer-id* must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword requirements which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer's context for possible parameter modification.

Values 1—32 egress

Values 1—8 egress

fc

Syntax **fc** *fc-name* [**create**]
no fc *fc-name*

Context config qos>queue-group-templates>egress>queue-group-template

Description The **fc** command is used to enter the forwarding class mapping context for the given fc-name. Each forwarding class has a default mapping depending on the egress queue group template. The system created

policer-output-queue template contains queues 1 and 2 by default with queue 1 being best-effort and queue 2 expedited. Forwarding classes be, l1, af and l2 all map to queue 1 by default. Forwarding classes h1, ef, h2 and nc all map to queue 2 by default. More queues may be created within the policer-output-queues template and the default forwarding classes may be changed to any defined queue within the template.

When all other user defined egress queue group templates are created, only queue 1 (best-effort) exists and all forwarding classes are mapped to that queue. Other queues may be created and the forwarding classes may be changed to any defined queue within the template.

Besides the default mappings within the templates, the egress queue group template forwarding class queue mappings operate the same as the forwarding class mappings in a sap-egress QoS policy.

The template forwarding class mappings are the default mechanism for mapping egress policed traffic to a queue within an egress port queue group associated with the template. If a queue-id is

explicitly specified in the QoS policy forwarding class policer mapping, and that queue exists within the queue group, the template forwarding class mapping is ignored. Egress policed subscriber traffic works in a slightly different way. The subscriber and subscriber host support destination and organization strings which when exist are used to identify the egress port queue group. In this instance, the forwarding class mappings are always used and any queue overrides in the QoS policy are ignored. If neither string exists for the subscriber host, the egress queue group queue-id can be derived from either the QoS policy policer mapping or the template forwarding class queue mappings.

The **no** form of this command is used to return the specified forwarding class to its default template queue mapping.

Parameters *fc-name* — A valid forwarding class must be specified as *fc-name* when the **fc** command is executed. When the **fc** *fc-name* command is successfully executed, the system will enter the specified forwarding class context where the **queue** *queue-id* command may be executed.

Values **be**, **l1**, **af**, **l2**, **h1**, **ef**, **h2** or **nc**

Default None

policer

Syntax **policer** *policer-id*
no policer

Context config>qos>queue-group-templates>ingress>queue-group
config>qos>queue-group-templates>egress>queue-group

Description This command is used in ingress and egress queue-group templates to create, modify or delete a policer.

Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The ingress queue-group template may have up to 32 policers (numbered 1 through 32) may be defined while the egress queue-group template supports a maximum of 8 (numbered 1 through 8). While a policer may be defined in a queue-group template, it is not actually created until the queue-group template is instantiated on ingress context of a forwarding plane or on the egress context of a port.

Once a policer is created, the policer's metering rate and profiling rates may be defined as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues which have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

Once a policer is created, it cannot be deleted from the queue-group template unless any forwarding classes that are redirected to the policer are first removed.

The **no** form of this command deletes the policer.

Parameters *policer-id* — The policer-id must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword requirements which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer's context for possible parameter modification.

Values 1—32 ingress

1—8 egress

port-redirect-group

Syntax **port-redirect-group {queue** *queue-id* **| policer** *policer-id* **[queue** *queue-id***]}**
no port-redirect-group

Context config>qos>network>egress>fc

Description This command is used to redirect the FC of a packet of a PW or network IP interface to an egress port queue-group.

It defines the mapping of a FC to a queue-id or a policer-id and a queue-id and redirects the lookup of the queue or policer of the same ID in some egress port queue-group instance. However, the queue-group name and instance are explicitly provided only at the time the network QoS policy is applied to egress context of a spoke-sdp or a network IP interface.

The no version of this command removes the redirection of the FC.

Parameters **queue** *queue-id* — The specified queue-id must exist within the queue-group template applied to the egress context of the port.

Values 1—8

policer *policer-id* — The specified policer-id must exist within the queue-group template applied to the egress context of the port

Values 1—8

fp-redirect-group

Syntax **fp-redirect-group policer** *policer-id*
no fp-redirect-group policer

Context config>qos>network>ingress>fc

Description This command is used to redirect the FC of a packet of a pseudowire or network IP interface to an ingress forwarding plane queue-group.

It defines the mapping of a FC to a policer-id and redirects the lookup of the policer of the same ID in some ingress forwarding plane queue-group instance. However, the queue-group name and instance are explicitly provided only at the time the network QoS policy is applied to the ingress context of a spoke-sdp or a network IP interface.

The **no** version of this command removes the redirection of the FC.

Parameters **policer** *policer-id* — The specified policer-id must exist within the queue-group template applied to the ingress context of the forwarding plane

Values 1—8

fp-redirect-group

Syntax **fp-redirect-group multicast-policer** *policer-id*
no fp-redirect-group multicast-policer

Context config>qos>network>ingress>fc

Description This command is used to redirect the FC of a multicast packet of a pseudowire or network IP interface to an ingress forwarding plane queue-group.

It defines the mapping of a FC to a policer-id and redirects the lookup of the policer of the same ID in some ingress forwarding plane queue-group instance. However, the queue-group name and instance are explicitly provided only at the time the network QoS policy is applied to the ingress context of a spoke-sdp or a network IP interface.

The no version of this command removes the redirection of the FC.

Parameters **multicast-policer** *policer-id* — The specified policer-id must exist within the queue-group template applied to the ingress context of the forwarding plane

Values 1—32

packet-byte-offset

Syntax **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}
no packet-byte-offset9

Context config>qos>queue-group-templates>egress>queue-group>queue>xp-specific

Description This command is used to modify the size of each packet handled by the queue by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed.

The packet-byte-offset command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers.

When a packet-byte-offset value is applied to a queue instance, it adjusts the immediate packet size. This means that the queue rates, i.e., operational PIR and CIR, and queue bucket updates use the adjusted packet size. In addition, the queue statistics will also reflect the adjusted packet size. Scheduler policy rates, which are data rates, will use the adjusted packet size.

The port scheduler max-rate and the priority level rates and weights, if a Weighted Scheduler Group is used, are always on-the-wire rates and thus use the actual frame size. The same goes for the agg-rate-limit on a SAP, a subscriber, or a Multi-Service Site (MSS) when the queue is port-parented.

When the user enables frame-based-accounting in a scheduler policy or queue-frame-based-accounting with agg-rate-limit in a port scheduler policy, the queue rate will be capped to a user configured on-the-wire rate but the packet-byte-offset value is still in effect as explained above.

The **no** version of this command is used to remove per packet size modifications from the queue.

Parameters **add** *bytes* — The add keyword is mutually exclusive to the subtract keyword. Either add or subtract must be specified. When add is defined, the corresponding bytes parameter specifies the number of bytes that is added to the size of each packet associated with the queue for scheduling and accounting purposes

Values 0—32

Default None

subtract *bytes* — The subtract keyword is mutually exclusive to the add keyword. Either subtract or add must be specified. When subtract is defined, the corresponding bytes parameter specifies the number of bytes that is subtracted to the size of each packet associated with the queue for scheduling and accounting purposes.

Values 0—64

Default None

queue

Syntax **queue** *queue-id* [*queue-type*] [**create**]
no queue *queue-id*

Context cfg>qos>qgrps>egr>qgrp

Description This command creates a queue for use in a queue group template. Once created, the defined queue-id acts as a repository for the default parameters for the queue. The template queue is created on each queue-group object which is created with the queue group template name. Each queue is identified within the template by a queue-id number. The template ensures that all queue groups created with the template name will have the same queue-ids providing a uniform structure for the forwarding class redirection commands in the SAP egress QoS policies. The parameters within the template queue will be used as the default settings for each queue in the actual queue group. The queue parameters may be individually changed for each queue in each queue group using per queue overrides.

The **no** form of the command

Default none

adaptation-rule

Syntax **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
no adaptation-rule

Context config>qos>qgrpid>egr>qgrp>queue
config>qos>qgrpid>ing>qgrp>queue
config>qos>qgrpid>ing>qgrp>policer

Description This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default adaptation-rule pir closest cir closest

Parameters **pir** — Defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — Defines the constraints enforced when adapting the CIR rate defined within the **queue** queue-id **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.

Values **max** — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

adv-config-policy

Syntax **adv-config-policy** *adv-config-policy-name*
no adv-config-policy

Context config>qos>qgrpid>egress>qgrp>policer
config>qos>qgrpid>ingress>qgrp>policer

Description This command specifies the name of the advanced configuration policy to be applied with this policer.

Parameters *adv-config-policy-name* — Specifies an existing advanced configuration policy up to 32 characters in length.

burst-limit

Syntax **burst-limit** {**default** | *size* [**byte** | **kilobyte**]}
no burst-limit

Context config>qos>qgrps>egr>qgrp>queue
config>qos>qgrpid>ing>qgrp>queue

Description The `queue burst-limit` command is used to define an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The `burst-limit` command is supported under the sap-ingress and sap-egress QoS policy queues. The command is also supported under the ingress and egress queue-group-templates queues.

The **no** form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies or queue group templates. When specified within a queue-override queue context, any current burst limit override for the queue will be removed and the queue's burst limit will be controlled by its defining policy or template.

Parameters **default —** The default parameter is mutually exclusive to specifying an explicit size value. When burst-limit default is executed, the queue is returned to the system default value.

size — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in Kilobytes. If the value is intended to be interpreted in bytes, the byte qualifier must be added following size.

Values 1 to 14,000 (14,000 or 14,000,000 depending on bytes or kilobytes)

Default No default for size, use the default keyword to specify default burst limit

byte — The **bytes** qualifier is used to specify that the value given for size must be interpreted as the burst limit in bytes. The byte qualifier is optional and mutually exclusive with the kilobytes qualifier.

kilobyte — The **kilobyte** qualifier is used to specify that the value given for size must be interpreted as the burst limit in Kilobytes. The kilobyte qualifier is optional and mutually exclusive with the bytes qualifier. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>qos>qgrps>egr>qgrp>queue
config>qos>qgrpid>ing>qgrp>queue
config>qos>qgrpid>egr>qgrp>policer
config>qos>qgrpid>ing>qgrp>policer

Description The **cbs** command is used to define the default committed buffer size for the template queue or the CBS for the template policer. Overall, the cbs command follows the same behavior and provisioning characteristics as the cbs command in the SAP ingress and egress QoS policy.

The **no** form of this command restores the default CBS size to the template policer.

Default default

Parameters *size-in-kbytes* — For the queues, the size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes). For policers, the size parameter is an integer expression of the number of kilobytes for the policer CBS.

Values Queues: 0 - 131072 or default

Values Policers:0 — 16777216 or **default**

high-prio-only

Syntax **high-prio-only** *percent*
no high-prio-only

Context config>qos>qgrps>egr>qgrp>queue
config>qos>qgrpid>ing>qgrp>queue
config>qos>qgrpid>ing>qgrp>policer

Description The **high-prio-only** command configures the percentage of buffer space for the queue, used exclusively by high priority packets. The specified value overrides the default value for the context.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

The **no** form of this command restores the default high priority reserved size.

Parameters *percent* — The percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10.

Values 0 — 100, default

mbs

Syntax **mbs** *size* [**bytes** | **kilobytes**]
no mbs

Context config>qos>qgrps>egr>qgrp>queue
config>qos>qgrpid>ing>qgrp>queue
config>qos>qgrpid>egr>qgrp>policer
config>qos>qgrpid>ing>qgrp>policer

Description For queues, the Maximum Burst Size (MBS) command the default maximum buffer size for the template queue. The value is given in kilobytes.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The port Ethernet access ingress and egress queue group context for mbs provides a mechanism for overriding the default maximum size for the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

For policers, this command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low priority violate threshold.

For ingress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold.

The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by high-prio-only is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.

The **no** form of this command returns the MBS size assigned to the queue to the value.

Default 64 kilobytes when PIR = max, otherwise 10ms volume of traffic for a configured non zero/non max PIR.

Parameters *size* [**bytes** | **kilobytes**] — For queues, the size parameter is an integer expression of the maximum number of bytes or kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets. For policers, the size parameter is an integer expression of the maximum number of bytes for the policer's MBS.

Values For queues: 0 — 134217728, default

Values For policers: 0 - 16777216, default

[**bytes** | **kilobytes**] — Select bytes or kilobytes. Kilobytes is the default.

parent

Syntax **parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
no parent

Context config>qos>qgrps>egr>qgrp>queue
config>qos>qgrpid>ing>qgrp>queue

Description This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the `config>qos>scheduler-policy>tier` *`level`* context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the

scheduler-name being directly or indirectly applied (through a multi-service customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The orphaned state must generate a log entry and a trap message. The SAP which the queue belongs to must also depict an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state mentioned above and automatically return to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of the command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters *scheduler-name* — The defined *scheduler-name* conforms to the same input criteria as the schedulers defined within a scheduler policy. Scheduler names are configured in the `config>qos>scheduler-policy>tier` *`level`* context. There are no checks performed at the time of definition to ensure that the *scheduler-name* exists within an existing scheduler policy. For the queue to use the defined *scheduler-name*, the scheduler exists on each egress SAP the queue is eventually created on. For the duration where *scheduler-name* does not exist on the egress SAP, the queue operates in an orphaned state.

Values Any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Default None. Each parental association must be explicitly defined.

weight *weight* — *weight* defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as weighted receive no parental bandwidth until all strict queues and schedulers on the parent have reached their maximum bandwidth or are idle. In this manner, weighted children are considered to be the lowest priority.

All **weight** values from all weighted active queues and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the queue or scheduler after the strict children are serviced. A weight is considered to be active when the pertaining queue or scheduler has not reached its maximum rate and still has packets to transmit. All child queues and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all strict and non-zero weighted queues and schedulers are operating at the maximum bandwidth or are idle.

Values 0 — 100

Default 1

level *level* — The optional **level** parameter defines the level of hierarchy when compared to other schedulers and queues when vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined

as **strict** receive no parental bandwidth until all strict queues and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

Children of the parent scheduler with a lower strict priority or that are weighted will not receive bandwidth until all children with a higher strict priority have either reached their maximum bandwidth or are idle. Children with the same strict level are serviced in a round robin fashion.

Values 1 — 100

Default 1

cir-weight *cir-weight* — Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 100

cir-level *cir-level* — Defines the port priority the queue or scheduler will use to receive bandwidth for its within-cir offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 8 (8 is the highest priority)

Default 0

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*]
no percent-rate

Context config>qos>queue-group-templates>egress>queue-group-template>queue

Description The **percent-rate** command within the egress queue group template enables support for a queue's PIR and CIR rate to be configured as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

The **no** form of this command returns the queue to its default shaping rate and cir rate.

Parameters *pir-percent* — This parameter is used to express the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation and the egress-rate setting.

Values Percentage ranging from 0.01 to 100.00. The default is 100.00.

cir-percent — The **cir** keyword is optional and when defined the required *pir-percent* parameter expresses the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation and the egress-rate setting.

Values Percentage ranging from 0.00 to 100.00. The default is 100.00.

port-parent

Syntax **port-parent** [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
no port-parent

Context config>qos>qgrps>egr>qgrp>queue

Description This command defines the port scheduling parameters used to control the queues behavior when a virtual egress port scheduling is enabled where the egress queue group template is applied. The port-parent command follows the same behavior and provisioning characteristics as the parent command in the SAP egress QoS policy. The port-parent command is mutually exclusive with the parent command.

The **no** form of the command removes the values from the configuration.

Default none

Parameters **weight *weight* —** Defines the weight the queue or scheduler will use at the above-cir port priority level (defined by the level parameter).

Values 0 — 100

Default 1

level *level* — Defines the port priority the queue or scheduler will use to receive bandwidth for its above-cir offered-load.

Values 1 — 8 (8 is the highest priority)

Default 1

cir-weight *cir-weight* — Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 100

cir-level *cir-level* — Defines the port priority the queue or scheduler will use to receive bandwidth for its within-cir offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or

scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 8 (8 is the highest priority)

Default 0

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>qos>qgrps>egr>qgrp>queue
config>qos>qgrpid>ing>qgrp>queue
config>qos>qgrps>ing>qgrp>policer
config>qos>qgrps>egr>qgrp>policer

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue or policer. The PIR defines the maximum rate that the queue or policer can transmit packets out an egress interface (for SAP egress queues or policer). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue or policer can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue or policer over other queues or policer competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue or policer's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id* or *policer-id*.

The **no** form of the command returns all queues or policer created with the *queue-id* or *policer-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default **rate max cir 0 —** The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue or policer. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to max. Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue or policer is provisioned.

Values Queue: 1— 100000000, **max**

Values Policer:1— 2000000000**, max**

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue or policer. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card then the CIR used is equivalent to max.

Values Queue: 0— 100000000, **max**

Values Policer:0— 2000000000, **max**

Default 0

xp-specific

Syntax **xp-specific**

Context config>qos>qgrps>egr>qgrp>queue

Description This command specifies queue parameters or behavior specific to the Q2 traffic management feature set. All IOMs within the XP family utilize the Q2 for traffic management queuing functions. When the SAP egress QoS policy is applied to a SAP on an IOM3-XP any commands and parameters defined within the xp-specific context will either override or augment the generic commands and parameters defined for the specific queue ID..

In the event that the QoS policy is applied to a SAP on a non-IOM3-XP, the commands and parameters within the xp-specific node are ignored.

When the QoS policy is applied to a LAG SAP that spans XP and non-XP IOMs, the **xp-specific** commands and parameters are applied for the SAP queues created on the IOM3-XP LAG links.

packet-byte-offset

Syntax **packet-byte-offset** {**add** *bytes* **| subtract** *bytes*}
no packet-byte-offset

Context config>qos>queue-group-templates>egress>queue-group>queue>xp-specific

Description This command is used to modify the size of each packet handled by the queue by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed.

The packet-byte-offset command is meant to be an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers.

When a packet-byte-offset value is applied to a queue instance, it adjusts the immediate packet size. This means that the queue rates (i.e., operational PIR and CIR) and queue bucket updates use the adjusted packet size. In addition, the queue statistics will also reflect the adjusted packet size. Scheduler policy rates, which are data rates, will use the adjusted packet size.

The port scheduler **max-rate** and the priority level rates and weights, if a Weighted Scheduler Group is used, are always on-the-wire rates and thus use the actual frame size. The same goes for the agg-rate-limit on a SAP, a subscriber, or a Multi-Service Site (MSS) when the queue is port-parented.

When the user enables **frame-based-accounting** in a scheduler policy or **queue-frame-based-accounting** with agg-rate-limit in a port scheduler policy, the queue rate will be capped to a user configured on-the-wire rate, but the packet-byte-offset value is still in effect as explained above.

The **no** version of this command is used to remove per packet size modifications from the queue.

Parameters **add** *bytes* — The **add** keyword is mutually exclusive to the **subtact** keyword. Either **add** or **subtract** must be specified. When **add** is defined, the corresponding bytes parameter specifies the number of bytes that is added to the size of each packet associated with the queue for scheduling and accounting purposes.

Values 0—32

subtract *bytes* — The **subtract** keyword is mututally exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When **subtract** is defined, the corresponding bytes parameter specifies the number of bytes that is subtracted to the size of each packet associated with the queue for scheduling and accounting purposes.

Values 0—64

wred-queue

Syntax **wred-queue** [**policy** *slope-policy-name*]
no wred-queue

Context config>qos>qgrps>egr>qgrp>queue>xp-specific

Description This command alters the generic buffer pool association of the queue for the purpose of allowing queue-specific WRED slopes with minimal provisioning. When the **wred-queue** command is defined and the queue ID is created on an IOM3-XP, a buffer pool is created specifically for the queue and the queue obtains all buffers from that pool. The size of the pool is the same as the size of the queue. In this manner, the WRED slopes that operate based on the pool's buffer utilization are also reacting to the congestion depth of the queue.

The size of the buffer pool is dictated by the queue's **mbs** parameter. The size of the reserved CBS portion of the buffer pool is dictated by the queue's **cbs** parameter. The provisioning characteristics of the **mbs** and **cbs** commands have not been changed.

In the case where the QoS policy is applied to a SAP on an IOM3-XP which has WRED queue support shut down (**config>card>fp>egress>wred-queue-control>shutdown**) the WRED buffer pool is created, but the queue will continue to map to either to its default pool or the pool defined in the **pool** command. If the **no shutdown** command is executed on the IOM, the queue will at that point be automatically moved to its own WRED pool.

Each pool created for a queue using the **wred-queue** command shares buffers with all other wred-queue-enabled queues on the same IOM3-XP. The WRED pool buffer management behavior is defined within the **config>card>fp>egress>wred-queue-control** context.

The WRED slopes within the pool are defined by the slope policy associated with the queue. When a policy is not explicitly defined, the default slope policy is used. The slope policy enables, disables and defines the relative geometry of the high and low WRED slopes in the pool. The policy also specifies the time average factor (TAF) used by the pool when calculating the weighted average pool depth.

As packets attempt to enter the egress queue, they are associated with either the high or low WRED slope based on the packets profile. If the packet is in-profile, the high slope is used. The low slope is used by out-of-profile packets. Each WRED slope performs a probability discard based on the current weighted average pool depth.

When wred-queue is enabled for a SAP egress queue on an IOM3-XP, the queue's **pool** and **hi-priority-only** commands are ignored.

The number of wred-queue-enabled queues allowed per IOM3-XP is hard coded to 7500.The **no** form of the command restores the generic buffer pool behavior to the queue. The WRED pool is removed from the system.

Parameters *slope-policy-name* — Overrides the default WRED slope policy with an explicit slope policy. The defined slope policy must exist or the command will fail.

queue

Syntax **queue** *queue-id*
no queue

Context config>qos>queue-group-templates>egress>queue-group-template>fc

Description This command is used to map the forwarding class to the specified *queue-id*. The specified *queue-id* must exist within the egress queue group template. Once a queue is defined in a forwarding class mapping, that queue cannot be deleted unless the forwarding class mapping is moved to another queue within the template. Other criteria may also exist preventing the queue from being deleted from the template such as an applied SAP egress QoS policy mapping to the queue.

Parameters *queue-id* — The specified *queue-id* must exist within the egress queue group template.

Values 1–8

Default Dependent on user or system created template.

ingress

Syntax **ingress**

Context config>qos>qgrps

Description This command enables the context to create ingress queue group templates. Ingress queue group templates can be applied to ingress ports to create an ingress queue group of the same name.

An ingress template must be created for a group-name prior to creating a queue group with the same name on an ingress port.

Default none

queue

Syntax **queue** *queue-id* [**multipoint**] [*queue-type*] [*queue-mode*] [**create**]
no queue *queue-id*

Context cfg>qos>qgrps>egr>qgrp
cfg>qos>qgrps>ing>qgrp

Description This command creates a queue for use in a queue group template. Once created, the defined queue-id acts as a repository for the default parameters for the queue. The template queue is created on each queue-group object which is created with the queue group template name. Each queue is identified within the template by a queue-id number. The template ensures that all queue groups created with the template? name will have the same queue-ids providing a uniform structure for the forwarding class redirection commands in the SAP ingress QoS policies. The parameters within the template queue will be used as the default settings for each queue in the actual queue group. The queue parameters may be individually changed for each queue in each queue group using per queue overrides.

Once a queue within a template is mapped by a forwarding class on any object, the queue may be edited, but not deleted.

The **no** form of the command removes a template queue from the queue group template. If the queue is specified as a forwarding class redirection target in any SAP ingress QoS policy, the command will fail.

Default none

Parameters *queue-id* — This required parameter identifies the queue that will either be created or edited within the queue group template.

Values 1 — 32

multipoint — This optional keyword creates an ingress multipoint queue. Multipoint queues in a queue group may be used by ingress VPLS for forwarding types multicast, broadcast or unknown within a forwarding class. For ingress IES and VPRN access SAPs, only multicast is supported. Multipoint queues are only supported on ingress queue group templates

queue-type — The queue types are mutually exclusive to each other.

Values **expedite** — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.
best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

queue-mode — These keywords are optional and mutually exclusive when creating a new template queue. The keywords specify how the queue manages ingress explicitly profiled packets.

Values **profile-mode** — Overrides the default priority mode of the queue and allows the adoption of color-aware profiling within the queue. Forwarding classes and sub-classes may be explicitly defined as in-profile or out-of-profile. Out-of-profile classified packets bypass the CIR rate associated with the queue reserving it for the undefined or in-profile classified packets. If the template queue is not defined as profile-mode and the packet redirected

to the queue is explicitly out-of-profile based on the classification rules, the queues within CIR bandwidth may be consumed by the packet.

priority-mode — Defines that the SAP ingress QoS policy priority classification result will be honored by the queue. Priority mode is the default mode of the queue. High priority packets are allowed into the queue up to the mbs size defined for the queue. Low priority packets are discarded at the low priority MBS threshold which is derived from applying the hi-prio-only percentage to the queues MBS and subtracting that result from the mbs size defined.

create — Keyword used to create the queue ID instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
rate *pir-rate* **police**
no rate

Context config>qos>qgrpid>ing>qgrp>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. For SAP ingress, the CIR also defines the rate that packets are considered in-profile by the system. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP ingress or SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default none

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 — 100000000, **max**

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer.

Values 0 — 100000000, **max**

Default 0

police — Specifies that the out of profile traffic feeding into the physical queue instance should be dropped. Using this keyword will override the bandwidth specified by the SAP ingress queue's administrative CIR.

If the **police** keyword is not specified, the individual queue group overrides may override both the defined shaping rate and the cir defined profiling rate. When police is defined, only the policing rate may be overridden.

policer

Syntax **policer** *policer-id* [**create**]
no policer

Context cfg>qos>qgrps>ing>qgrp

Description This command configures a QoS ingress queue-group policer.

Default none

Parameters *policer-id* — This required parameter identifies the queue-group policer that will either be created or edited within the queue group template.

Values 1 — 32

create — This optional keyword creates an ingress queue-group policer.

profile-capped

Syntax **profile-capped**
no profile-capped

Context cfg>qos>qgrps>ing>qgrp>policer

Description This command enables a limit on the profile.

Default no profile-capped

packet-byte-offset

Syntax **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}
no packet-byte-offset

Context cfg>qos>qgrps>ing>qgrp>policer

Description This command configures a packet byte offset for the QoS ingress queue-group policer.

Default none

Parameters **add** *bytes* — Specifies an number of bytes to add as the offset amount.

Values 0 — 31

subtract *bytes* — Specifies an number of bytes to add as the offset amount.

Values 1 — 32

parent

Syntax **parent** {**root |** *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]
no parent

Context config>qos>qgrpid>ing>qgrp>policer

Description This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

The **no** form of the command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters **root —** Selects the root level arbiter for the parent to the child.

arbiter-name — Specifies an arbiter for the parent to the child.

Values Any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Default None. Each parental association must be explicitly defined.

weight-within-level — weight defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as weighted receive no parental bandwidth until all strict queues and schedulers on the parent have reached their maximum bandwidth or are idle. In this manner, weighted children are considered to be the lowest priority.

All **weight** values from all weighted active queues and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the queue or scheduler after the strict children are serviced. A weight is considered to be active when the pertaining queue or scheduler has not reached its maximum rate and still has packets to transmit. All child queues and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all strict and non-zero weighted queues and schedulers are operating at the maximum bandwidth or are idle.

Values 0 — 100

Default 1

level — The optional **level** parameter defines the level of hierarchy when compared to other schedulers and queues when vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as **strict** receive no parental bandwidth until all strict queues and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

Children of the parent scheduler with a lower strict priority or that are weighted will not receive bandwidth until all children with a higher strict priority have either reached their maximum bandwidth or are idle. Children with the same strict level are serviced in a round robin fashion.

Values 1 — 8

Default 1

stat-mode

Syntax **stat-mode** {**no-stats | minimal | offered-profile-no-cir | offered-priority-no-cir | offered-limited-profile-cir | offered-profile-cir | offered-priority-cir | offered-total-cir | offered-profile-capped-cir | offered-limited-capped-cir**}

Context cfg>qos>qgrps>ing>qgrp>policer

Description This command selects the statistics mode for the QoS ingress queue-group policer.

Default none

Parameters **no-stat —** Selects no statistics as the statistics mode.

minimal — Selects minimal statistics as the statistics mode.

offered-profile-no-cir — Selects no offered profile CIR statistics as the statistics mode.

offered-priority-no-cir — Selects no offered priority CIR statistics as the statistics mode.

offered-limited-profile-cir — Selects limited profile CIR statistics as the statistics mode.

offered-profile-cir — Selects offered profile CIR statistics as the statistics mode.

offered-priority-cir — Selects offered priority CIR statistics as the statistics mode.

offered-total-cir — Selects total statistics as the statistics mode.

offered-profile-capped-cir — Selects offered profile capped statistics as the statistics mode.

offered-limited-capped-cir — Selects offered limited capped statistics as the statistics mode.

Show Commands

queue-group

Syntax **queue-group** [*queue-group-name*] [**ingress** | **egress**] [**association** | **detail | summary**]
queue-group summary

Context show>qos

Description This command displays queue-group information.

Parameters *queue-group-name* — Specifies the name of an existing queue group template up to 32 characters in length.

ingress — Specifies whether the queue group name is an ingress policy.

egress — Specifies whether the queue group name is an egress policy.

associations — Displays the entities associated with the specified queue group name.

detail — Displays detailed queue group information for the specified queue group name.

summary — Displays the total number of queue-group instance per card (IOM).

Sample Output

```
*A:Dut-T>cfg>qos>qgrps>egr>qgrp# show qos queue-group egress
===============================================================================
Queue Group Egress
===============================================================================
Group-Name                     Description
-------------------------------------------------------------------------------
QG_egress_1                    Egress queue group
===============================================================================
*A:Dut-T#


*A:Dut-T# show qos queue-group egress QG_egress_1 detail
===============================================================================
QoS Queue-Group Egress
===============================================================================
-------------------------------------------------------------------------------
QoS Queue Group
-------------------------------------------------------------------------------
Group-Name     : QG_egress_1
Description    : Egress queue group
-------------------------------------------------------------------------------
Queue CIR Admin PIR Admin CBS     HiPrio PIR Lvl/Wt    Parent
      CIR Rule  PIR Rule  MBS             CIR Lvl/Wt
      Named-Buffer Pool
-------------------------------------------------------------------------------
1     0         max       def     def    1/1           None
      closest   closest   100             0/1
      (not-assigned)
2     0         max       def     def    1/1           None
```

```
      closest   closest   100            0/1
      (not-assigned)
3     0         max       def     def    1/1           None
      closest   closest   100            0/1
      (not-assigned)
4     0         max       def     def    1/1           None
      closest   closest   100            0/1
      (not-assigned)
===============================================================================
Queue Group Ports (access)
===============================================================================
Port                Sched Pol           Acctg Pol Stats     Description
-------------------------------------------------------------------------------
9/2/1                                   0         No
9/2/2                                   0         No
-------------------------------------------------------------------------------
===============================================================================
Queue Group Ports (network)
===============================================================================
Port                Sched Pol           Acctg Pol Stats     Description
-------------------------------------------------------------------------------
6/1/1                                   0         No
-------------------------------------------------------------------------------
===============================================================================
Queue Group Sap FC Maps
===============================================================================
Sap Policy     FC Name             Queue Id
-------------------------------------------------------------------------------
10             af                  2
10             be                  1
10             ef                  3
10             nc                  4
-------------------------------------------------------------------------------
Entries found: 4
-------------------------------------------------------------------------------
===============================================================================
*A:Dut-T#


*A:Dut-T# show qos queue-group egress QG_egress_1 association
===============================================================================
QoS Queue-Group Egress
===============================================================================
-------------------------------------------------------------------------------
QoS Queue Group
-------------------------------------------------------------------------------
Group-Name     : QG_egress_1
Description    : Egress queue group
===============================================================================
Queue Group Ports (access)
===============================================================================
Port                Sched Pol           Acctg Pol Stats     Description
-------------------------------------------------------------------------------
9/2/1                                   0         No
9/2/2                                   0         No
-------------------------------------------------------------------------------
===============================================================================
Queue Group Ports (network)
===============================================================================
```

```
Port                Sched Pol           Acctg Pol Stats     Description
-------------------------------------------------------------------------------
6/1/1                                   0         No
-------------------------------------------------------------------------------
===============================================================================
Queue Group Sap FC Maps
===============================================================================
Sap Policy     FC Name             Queue Id
-------------------------------------------------------------------------------
10             af                  2
10             be                  1
10             ef                  3
10             nc                  4
-------------------------------------------------------------------------------
Entries found: 4
-------------------------------------------------------------------------------
===============================================================================
*A:Dut-T#
*A:Dut-T# show qos queue-group summary
============================================================
card | access-ingress | access-egress | network-egress

   1 |        60      |       2047    |       0
   2 |        60      |         0     |      2047
============================================================
Total ingress QG templates per system: <num>
Total egress QG templates per system:  <num>
```

The total number of queue-group instance per card (IOM).

```
*A:Dut-T# show qos queue-group ingress
===============================================================================
Queue Group Ingress
===============================================================================
Group-Name                    Description
-------------------------------------------------------------------------------
QG_ingress_1                  Ingress queue-group
===============================================================================
*A:Dut-T#


*A:Dut-T# show qos queue-group ingress detail
===============================================================================
QoS Queue-Group Ingress
===============================================================================
-------------------------------------------------------------------------------
QoS Queue Group
-------------------------------------------------------------------------------
Group-Name     : QG_ingress_1
Description    : Ingress queue-group
-------------------------------------------------------------------------------
Queue Mode    CIR Admin PIR Admin CBS     HiPrio  PIR Lvl/Wt    Parent
              CIR Rule  PIR Rule  MBS             CIR Lvl/Wt
              Named-Buffer Pool
-------------------------------------------------------------------------------
1     Prio      0         max     def    def           1/1            None
                closest   closest 100                  0/1
```

```
                (not-assigned)
2     Prio      0         max     def    def           1/1            None
                closest   closest 100                  0/1
                (not-assigned)
3     Prio      0         max     def    def           1/1            None
                closest   closest 100                  0/1
                (not-assigned)
4     Prio      0         max     def    def           1/1            None
                closest   closest 100                  0/1
                (not-assigned)

===============================================================================
Queue Group Ports
===============================================================================
Port                Sched Pol           Acctg Pol Stats     Description
-------------------------------------------------------------------------------
9/2/1                                   0         No
9/2/2                                   0         No
-------------------------------------------------------------------------------
===============================================================================
Queue Group Sap FC Maps
===============================================================================
Sap Policy     FC Name             Queue (id type)
-------------------------------------------------------------------------------
10             af                    (2 Unicast)
10             be                    (1 Unicast)
10             ef                    (3 Unicast)
10             nc                    (4 Unicast)
-------------------------------------------------------------------------------
Entries found: 4
-------------------------------------------------------------------------------
===============================================================================
*A:Dut-T#


*A:Dut-T# show qos queue-group ingress association
===============================================================================
QoS Queue-Group Ingress
===============================================================================
-------------------------------------------------------------------------------
QoS Queue Group
-------------------------------------------------------------------------------
Group-Name     : QG_ingress_1
Description    : Ingress queue-group
===============================================================================
Queue Group Ports
===============================================================================
Port                Sched Pol           Acctg Pol Stats     Description
-------------------------------------------------------------------------------
9/2/1                                   0         No
9/2/2                                   0         No
-------------------------------------------------------------------------------
===============================================================================
Queue Group Sap FC Maps
===============================================================================
Sap Policy     FC Name             Queue (id type)
-------------------------------------------------------------------------------
10             af                    (2 Unicast)
10             be                    (1 Unicast)
```

```
10             ef                   (3 Unicast)
10             nc                   (4 Unicast)
-------------------------------------------------------------------------------
Entries found: 4
-------------------------------------------------------------------------------
===============================================================================
*A:Dut-T#


*A:Dut-T# show qos queue-group summary
===============================================================================
Queue-group instances per card
===============================================================================
card      port-acc-ing  port-acc-egr  port-nw-egr   fp-acc-ing    fp-nw-ing
-------------------------------------------------------------------------------
1         0             0             0             0             0
2         0             0             0             0             0
3         0             0             0             0             0
4         0             2             1000          0             500
5         0             0             0             0             0
6         0             0             0             0             0
7         0             0             0             0             0
8         0             0             0             0             0
9         0             2             1000          0             500
10        0             0             0             0             0
11        0             0             0             0             0
12        0             0             0             0             0
-------------------------------------------------------------------------------
Total ingress QG templates per system :  3
Total egress  QG templates per system :  5
===============================================================================
*A:Dut-T#
```

Related queue-group command output:

```
*A:Dut-T# show card 9 fp 1 ingress queue-group "QGIng1" mode network instance 1 statistics
===============================================================================
Card:9  Net.QGrp: QGIng1  Instance: 1
===============================================================================
Group Name    : QGIng1
Description   : (Not Specified)
Pol Ctl Pol   : pcp                     Acct Pol      : None
Collect Stats : disabled
-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  1  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         91836202                91465530792
Dro. All              :         6678807                 6649127172
For. All              :         85157395                84816403620

Ing. Policer:  2  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         93584703                90933906888
Dro. All              :         8320200                 6106644900
For. All              :         85264503                84827261988

Ing. Policer:  3  Grp: QGIng1 (Stats mode: minimal)
```

```
Off. All              :         93584703                90933906888
Dro. All              :         8320049                 6106288404
For. All              :         85264654                84827618484

Ing. Policer:  4  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         93584703                90933906888
Dro. All              :         8326509                 6110568864
For. All              :         85258194                84823338024

Ing. Policer:  5  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         93584703                90933906888
Dro. All              :         24877143                22616873028
For. All              :         68707560                68317033860

Ing. Policer:  6  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         93434643                90919501128
Dro. All              :         24727111                22602499656
For. All              :         68707532                68317001472

Ing. Policer:  7  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         93584703                90933906888
Dro. All              :         24877214                22616941944
For. All              :         68707489                68316964944

Ing. Policer:  8  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         93430663                90919119048
Dro. All              :         24723280                22602263280
For. All              :         68707383                68316855768

Ing. Policer:  9  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         0                       0
Dro. All              :         0                       0
For. All              :         0                       0

Ing. Policer: 10  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         0                       0
Dro. All              :         0                       0
For. All              :         0                       0

Ing. Policer: 11  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         0                       0
Dro. All              :         0                       0
For. All              :         0                       0

Ing. Policer: 12  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         0                       0
Dro. All              :         0                       0
For. All              :         0                       0

Ing. Policer: 13  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         0                       0
Dro. All              :         0                       0
For. All              :         0                       0

Ing. Policer: 14  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         0                       0
Dro. All              :         0                       0
For. All              :         0                       0
```

```
Ing. Policer: 15  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         0                       0
Dro. All              :         0                       0
For. All              :         0                       0

Ing. Policer: 16  Grp: QGIng1 (Stats mode: minimal)
Off. All              :         0                       0
Dro. All              :         0                       0
For. All              :         0                       0
===============================================================================
*A:Dut-T#


*A:Dut-T# show qos policer-hierarchy card 9 fp 1 queue-group "QGIng1" ingress instance 1
detail
===============================================================================
Policer Hierarchy - Card: 9 Queue-Group: QGIng1
===============================================================================
Ingress Policer Policy         :
-------------------------------------------------------------------------------
Legend :
(*) real-time dynamic value
(w) Wire rates
-------------------------------------------------------------------------------
root (Ing)
|
No Active Access Members Found on slot 9

|
| slot(9) (Network)
|    Profile-preferred:Disabled
|    MaxPIR:1500
|    ConsumedByChildren:1500
|    OperPIR:1500        OperFIR:1500
|
|    DepthPIR:205904 bytes
|  Priority 8
|    Oper Thresh Unfair:311296      Oper Thresh Fair:425984
|    Association count:0
|  Priority 7
|    Oper Thresh Unfair:311296      Oper Thresh Fair:425984
|    Association count:0
|  Priority 6
|    Oper Thresh Unfair:311296      Oper Thresh Fair:425984
|    Association count:0
|  Priority 5
|    Oper Thresh Unfair:311296      Oper Thresh Fair:425984
|    Association count:0
|  Priority 4
|    Oper Thresh Unfair:311296      Oper Thresh Fair:425984
|    Association count:0
|  Priority 3
|    Oper Thresh Unfair:311296      Oper Thresh Fair:425984
|    Association count:0
|  Priority 2
|    Oper Thresh Unfair:311296      Oper Thresh Fair:425984
|    Association count:4
|  Priority 1
|    Oper Thresh Unfair:102400      Oper Thresh Fair:204800
```

```
|    Association count:4
|
|--(A) : DATA (QGrp: QGIng1 Instance: 1 )
|   |    MaxPIR:max
|   |    ConsumedByChildren:500
|   |    OperPIR:500         OperFIR:500
|   |
|   |    [Level 1 Weight 1]
|   |    Assigned PIR:500        Offered:41603
|   |    Consumed:500
|   |
|   |    Assigned FIR:500
|   |
|   |--(P) : Policer Net-FPQG-1-T3:1->4
|   |   |    MaxPIR:max          MaxCIR:0
|   |   |    CBS:0               MBS:20480
|   |   |    HiPrio:2048
|   |   |    Depth:1184
|   |   |
|   |   |    OperPIR:128          OperCIR:0
|   |   |    OperFIR:128
|   |   |    PacketByteOffset:0
|   |   |    StatMode: minimal
|   |   |
|   |   |    [Level 1 Weight 1]
|   |   |    Assigned PIR:125        Offered:9966
|   |   |    Consumed:125
|   |   |
|   |   |    Assigned FIR:125
|   |
|   |--(P) : Policer Net-FPQG-1-T3:1->3
|   |   |    MaxPIR:max          MaxCIR:0
|   |   |    CBS:0               MBS:20480
|   |   |    HiPrio:2048
|   |   |    Depth:18256
|   |   |
|   |   |    OperPIR:128          OperCIR:0
|   |   |    OperFIR:128
|   |   |    PacketByteOffset:0
|   |   |    StatMode: minimal
|   |   |
|   |   |    [Level 1 Weight 1]
|   |   |    Assigned PIR:125        Offered:9966
|   |   |    Consumed:125
|   |   |
|   |   |    Assigned FIR:125
|   |
|   |--(P) : Policer Net-FPQG-1-T3:1->2
|   |   |    MaxPIR:max          MaxCIR:0
|   |   |    CBS:0               MBS:20480
|   |   |    HiPrio:2048
|   |   |    Depth:18944
|   |   |
|   |   |    OperPIR:128          OperCIR:0
|   |   |    OperFIR:128
|   |   |    PacketByteOffset:0
|   |   |    StatMode: minimal
|   |   |
|   |   |    [Level 1 Weight 1]
```

```
|   |   |    Assigned PIR:125        Offered:9967
|   |   |    Consumed:125
|   |   |
|   |   |    Assigned FIR:125
|   |
|   |--(P) : Policer Net-FPQG-1-T3:1->1
|   |   |    MaxPIR:max          MaxCIR:0
|   |   |    CBS:0               MBS:20480
|   |   |    HiPrio:2048
|   |   |    Depth:19024
|   |   |
|   |   |    OperPIR:128          OperCIR:0
|   |   |    OperFIR:128
|   |   |    PacketByteOffset:0
|   |   |    StatMode: minimal
|   |   |
|   |   |    [Level 1 Weight 1]
|   |   |    Assigned PIR:125        Offered:11724
|   |   |    Consumed:125
|   |   |
|   |   |    Assigned FIR:125
|
|--(A) : HIGH (QGrp: QGIng1 Instance: 1 )
|   |    MaxPIR:1000
|   |    ConsumedByChildren:1000
|   |    OperPIR:1000        OperFIR:1000
|   |
|   |    [Level 2 Weight 1]
|   |    Assigned PIR:1000       Offered:1000
|   |    Consumed:1000
|   |
|   |    Assigned FIR:1000
|   |
|   |--(P) : Policer Net-FPQG-1-T3:1->8
|   |   |    MaxPIR:max          MaxCIR:0
|   |   |    CBS:0               MBS:20480
|   |   |    HiPrio:2048
|   |   |    Depth:21353
|   |   |
|   |   |    OperPIR:250          OperCIR:0
|   |   |    OperFIR:250
|   |   |    PacketByteOffset:0
|   |   |    StatMode: minimal
|   |   |
|   |   |    [Level 1 Weight 1]
|   |   |    Assigned PIR:250        Offered:9966
|   |   |    Consumed:250
|   |   |
|   |   |    Assigned FIR:250
|   |
|   |--(P) : Policer Net-FPQG-1-T3:1->7
|   |   |    MaxPIR:max          MaxCIR:0
|   |   |    CBS:0               MBS:20480
|   |   |    HiPrio:2048
|   |   |    Depth:21065
|   |   |
|   |   |    OperPIR:250          OperCIR:0
|   |   |    OperFIR:250
|   |   |    PacketByteOffset:0
```

```
|   |   |    StatMode: minimal
|   |   |
|   |   |    [Level 1 Weight 1]
|   |   |    Assigned PIR:250        Offered:9967
|   |   |    Consumed:250
|   |   |
|   |   |    Assigned FIR:250
|   |
|   |--(P) : Policer Net-FPQG-1-T3:1->6
|   |   |    MaxPIR:max          MaxCIR:0
|   |   |    CBS:0               MBS:20480
|   |   |    HiPrio:2048
|   |   |    Depth:21353
|   |   |
|   |   |    OperPIR:250          OperCIR:0
|   |   |    OperFIR:250
|   |   |    PacketByteOffset:0
|   |   |    StatMode: minimal
|   |   |
|   |   |    [Level 1 Weight 1]
|   |   |    Assigned PIR:250        Offered:9967
|   |   |    Consumed:250
|   |   |
|   |   |    Assigned FIR:250
|   |
|   |--(P) : Policer Net-FPQG-1-T3:1->5
|   |   |    MaxPIR:max          MaxCIR:0
|   |   |    CBS:0               MBS:20480
|   |   |    HiPrio:2048
|   |   |    Depth:21065
|   |   |
|   |   |    OperPIR:250          OperCIR:0
|   |   |    OperFIR:250
|   |   |    PacketByteOffset:0
|   |   |    StatMode: minimal
|   |   |
|   |   |    [Level 1 Weight 1]
|   |   |    Assigned PIR:250        Offered:9967
|   |   |    Consumed:250
|   |   |
|   |   |    Assigned FIR:250


===============================================================================
*A:Dut-T#


*A:Dut-T# show qos policer port 9/2/4 network egress queue-group "QGEgr1" instance 1
===============================================================================
Policer Information (Summary), Slot 9
===============================================================================
-------------------------------------------------------------------------------
Name                FC-Maps       MBS       HP-Only A.PIR    A.CIR
Direction                         CBS       Depth   O.PIR    O.CIR    O.FIR
-------------------------------------------------------------------------------
Net-PQG-9/2/4-QGEgr1:1->8
Egress                            64 KB     8 KB    Max      0
                                  0 KB      1026    Max      0        Max
Net-PQG-9/2/4-QGEgr1:1->7
```

```
Egress                          64 KB     8 KB    Max      0
                                0 KB      1026    Max      0        Max
Net-PQG-9/2/4-QGEgr1:1->6
Egress                          64 KB     8 KB    Max      0
                                0 KB      1026    Max      0        Max
Net-PQG-9/2/4-QGEgr1:1->5
Egress                          64 KB     8 KB    Max      0
                                0 KB      1026    Max      0        Max
Net-PQG-9/2/4-QGEgr1:1->4
Egress                          64 KB     8 KB    Max      0
                                0 KB      1026    Max      0        Max
Net-PQG-9/2/4-QGEgr1:1->3
Egress                          64 KB     8 KB    Max      0
                                0 KB      1026    Max      0        Max
Net-PQG-9/2/4-QGEgr1:1->2
Egress                          64 KB     8 KB    Max      0
                                0 KB      1026    Max      0        Max
Net-PQG-9/2/4-QGEgr1:1->1
Egress                          64 KB     8 KB    Max      0
                                0 KB      1026    Max      0        Max
===============================================================================
*A:Dut-T#


*A:Dut-T# show qos policer port 9/2/4 network egress queue-group "QGEgr1" instance 1 detail
===============================================================================
Policer Info (Net-PQG-9/2/4-QGEgr1:1->8), Slot 9
===============================================================================
Policer Name       : Net-PQG-9/2/4-QGEgr1:1->8
Direction          : Egress             Fwding Plane       : 1
Depth PIR          : 1026 Bytes         Depth CIR          : 0 Bytes
Depth FIR          : 1026 Bytes
MBS                : 64 KB              CBS                : 0 KB
Hi Prio Only       : 8 KB               Pkt Byte Offset    : 0
Admin PIR          : Max                Admin CIR          : 0 Kbps
Oper PIR           : Max                Oper CIR           : 0 Kbps
Oper FIR           : Max
Stat Mode          : minimal
Parent Arbiter Name: (Not Specified)
-------------------------------------------------------------------------------
Arbiter Member Information
-------------------------------------------------------------------------------
Offered Rate       : 0 Kbps
Level              : 0                  Weight             : 0
Parent PIR         : 0 Kbps             Parent FIR         : 0 Kbps
Consumed           : 0 Kbps
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Policer Info (Net-PQG-9/2/4-QGEgr1:1->7), Slot 9
===============================================================================
Policer Name       : Net-PQG-9/2/4-QGEgr1:1->7
Direction          : Egress             Fwding Plane       : 1
Depth PIR          : 1026 Bytes         Depth CIR          : 0 Bytes
Depth FIR          : 1026 Bytes
MBS                : 64 KB              CBS                : 0 KB
Hi Prio Only       : 8 KB               Pkt Byte Offset    : 0
Admin PIR          : Max                Admin CIR          : 0 Kbps
Oper PIR           : Max                Oper CIR           : 0 Kbps
```

```
Oper FIR           : Max
Stat Mode          : minimal
Parent Arbiter Name: (Not Specified)
-------------------------------------------------------------------------------
Arbiter Member Information
-------------------------------------------------------------------------------
Offered Rate       : 0 Kbps
Level              : 0                 Weight             : 0
Parent PIR         : 0 Kbps            Parent FIR         : 0 Kbps
Consumed           : 0 Kbps
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Policer Info (Net-PQG-9/2/4-QGEgr1:1->6), Slot 9
===============================================================================
Policer Name       : Net-PQG-9/2/4-QGEgr1:1->6
Direction          : Egress            Fwding Plane       : 1
Depth PIR          : 1026 Bytes        Depth CIR          : 0 Bytes
Depth FIR          : 1026 Bytes
MBS                : 64 KB             CBS                : 0 KB
Hi Prio Only       : 8 KB              Pkt Byte Offset    : 0
Admin PIR          : Max               Admin CIR          : 0 Kbps
Oper PIR           : Max               Oper CIR           : 0 Kbps
Oper FIR           : Max
Stat Mode          : minimal
Parent Arbiter Name: (Not Specified)
-------------------------------------------------------------------------------
Arbiter Member Information
-------------------------------------------------------------------------------
Offered Rate       : 0 Kbps
Level              : 0                 Weight             : 0
Parent PIR         : 0 Kbps            Parent FIR         : 0 Kbps
Consumed           : 0 Kbps
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Policer Info (Net-PQG-9/2/4-QGEgr1:1->5), Slot 9
===============================================================================
Policer Name       : Net-PQG-9/2/4-QGEgr1:1->5
Direction          : Egress            Fwding Plane       : 1
Depth PIR          : 1026 Bytes        Depth CIR          : 0 Bytes
Depth FIR          : 1026 Bytes
MBS                : 64 KB             CBS                : 0 KB
Hi Prio Only       : 8 KB              Pkt Byte Offset    : 0
Admin PIR          : Max               Admin CIR          : 0 Kbps
Oper PIR           : Max               Oper CIR           : 0 Kbps
Oper FIR           : Max
Stat Mode          : minimal
Parent Arbiter Name: (Not Specified)
-------------------------------------------------------------------------------
Arbiter Member Information
-------------------------------------------------------------------------------
Offered Rate       : 0 Kbps
Level              : 0                 Weight             : 0
Parent PIR         : 0 Kbps            Parent FIR         : 0 Kbps
Consumed           : 0 Kbps
-------------------------------------------------------------------------------
===============================================================================
```

```
===============================================================================
Policer Info (Net-PQG-9/2/4-QGEgr1:1->4), Slot 9
===============================================================================
Policer Name       : Net-PQG-9/2/4-QGEgr1:1->4
Direction          : Egress             Fwding Plane       : 1
Depth PIR          : 1026 Bytes         Depth CIR          : 0 Bytes
Depth FIR          : 1026 Bytes
MBS                : 64 KB              CBS                : 0 KB
Hi Prio Only       : 8 KB               Pkt Byte Offset    : 0
Admin PIR          : Max                Admin CIR          : 0 Kbps
Oper PIR           : Max                Oper CIR           : 0 Kbps
Oper FIR           : Max
Stat Mode          : minimal
Parent Arbiter Name: (Not Specified)
-------------------------------------------------------------------------------
Arbiter Member Information
-------------------------------------------------------------------------------
Offered Rate       : 0 Kbps
Level              : 0                  Weight             : 0
Parent PIR         : 0 Kbps             Parent FIR         : 0 Kbps
Consumed           : 0 Kbps
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Policer Info (Net-PQG-9/2/4-QGEgr1:1->3), Slot 9
===============================================================================
Policer Name       : Net-PQG-9/2/4-QGEgr1:1->3
Direction          : Egress             Fwding Plane       : 1
Depth PIR          : 1026 Bytes         Depth CIR          : 0 Bytes
Depth FIR          : 1026 Bytes
MBS                : 64 KB              CBS                : 0 KB
Hi Prio Only       : 8 KB               Pkt Byte Offset    : 0
Admin PIR          : Max                Admin CIR          : 0 Kbps
Oper PIR           : Max                Oper CIR           : 0 Kbps
Oper FIR           : Max
Stat Mode          : minimal
Parent Arbiter Name: (Not Specified)
-------------------------------------------------------------------------------
Arbiter Member Information
-------------------------------------------------------------------------------
Offered Rate       : 0 Kbps
Level              : 0                  Weight             : 0
Parent PIR         : 0 Kbps             Parent FIR         : 0 Kbps
Consumed           : 0 Kbps
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Policer Info (Net-PQG-9/2/4-QGEgr1:1->2), Slot 9
===============================================================================
Policer Name       : Net-PQG-9/2/4-QGEgr1:1->2
Direction          : Egress             Fwding Plane       : 1
Depth PIR          : 1026 Bytes         Depth CIR          : 0 Bytes
Depth FIR          : 1026 Bytes
MBS                : 64 KB              CBS                : 0 KB
Hi Prio Only       : 8 KB               Pkt Byte Offset    : 0
Admin PIR          : Max                Admin CIR          : 0 Kbps
Oper PIR           : Max                Oper CIR           : 0 Kbps
Oper FIR           : Max
Stat Mode          : minimal
```

```
Parent Arbiter Name: (Not Specified)
-------------------------------------------------------------------------------
Arbiter Member Information
-------------------------------------------------------------------------------
Offered Rate       : 0 Kbps
Level              : 0                  Weight             : 0
Parent PIR         : 0 Kbps             Parent FIR         : 0 Kbps
Consumed           : 0 Kbps
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Policer Info (Net-PQG-9/2/4-QGEgr1:1->1), Slot 9
===============================================================================
Policer Name       : Net-PQG-9/2/4-QGEgr1:1->1
Direction          : Egress             Fwding Plane       : 1
Depth PIR          : 1026 Bytes         Depth CIR          : 0 Bytes
Depth FIR          : 1026 Bytes
MBS                : 64 KB              CBS                : 0 KB
Hi Prio Only       : 8 KB               Pkt Byte Offset    : 0
Admin PIR          : Max                Admin CIR          : 0 Kbps
Oper PIR           : Max                Oper CIR           : 0 Kbps
Oper FIR           : Max
Stat Mode          : minimal
Parent Arbiter Name: (Not Specified)
-------------------------------------------------------------------------------
Arbiter Member Information
-------------------------------------------------------------------------------
Offered Rate       : 0 Kbps
Level              : 0                  Weight             : 0
Parent PIR         : 0 Kbps             Parent FIR         : 0 Kbps
Consumed           : 0 Kbps
-------------------------------------------------------------------------------
===============================================================================
-------------------------------------------------------------------------------
Network Interface Association
-------------------------------------------------------------------------------
No Association Found.
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
SDP Association
-------------------------------------------------------------------------------
Policer Info (1->1:101->10), Slot 9
Policer Info (1->2:102->10), Slot 9
Policer Info (1->3:103->10), Slot 9
Policer Info (1->4:104->10), Slot 9
Policer Info (1->5:105->10), Slot 9
Policer Info (1->6:106->10), Slot 9
Policer Info (1->7:107->10), Slot 9
Policer Info (1->8:108->10), Slot 9
Policer Info (1->9:109->10), Slot 9
Policer Info (1->10:110->10), Slot 9
Policer Info (1->11:111->10), Slot 9
Policer Info (1->12:112->10), Slot 9
Policer Info (1->13:113->10), Slot 9
Policer Info (1->14:114->10), Slot 9
Policer Info (1->15:115->10), Slot 9
Policer Info (1->16:116->10), Slot 9
```

```
*A:Dut-T# show port 9/2/4 queue-group egress "QGEgr1" statistics instance 1
-------------------------------------------------------------------------------
Ethernet port 9/2/4 Network Egress queue-group
-------------------------------------------------------------------------------
                        Packets                 Octets

Egress Queue:  1    Group: QGEgr1    Instance-Id:  1
In Profile forwarded  : 0                       0
In Profile dropped    : 0                       0
Out Profile forwarded : 0                       0
Out Profile dropped   : 0                       0

Egress Policer:  1  Group: QGEgr1  Instance-Id: 1
Stats mode: minimal
Off. All              : 22133354                22708821204
Dro. All              : 0                       0
For. All              : 22133354                22708821204

Egress Policer:  2  Group: QGEgr1  Instance-Id: 1
Stats mode: minimal
Off. All              : 22133354                22708821204
Dro. All              : 0                       0
For. All              : 22133354                22708821204

Egress Policer:  3  Group: QGEgr1  Instance-Id: 1
Stats mode: minimal
Off. All              : 22133354                22708821204
Dro. All              : 0                       0
For. All              : 22133354                22708821204

Egress Policer:  4  Group: QGEgr1  Instance-Id: 1
Stats mode: minimal
Off. All              : 22133354                22708821204
Dro. All              : 0                       0
For. All              : 22133354                22708821204

Egress Policer:  5  Group: QGEgr1  Instance-Id: 1
Stats mode: minimal
Off. All              : 22133355                22708822230
Dro. All              : 0                       0
For. All              : 22133355                22708822230

Egress Policer:  6  Group: QGEgr1  Instance-Id: 1
Stats mode: minimal
Off. All              : 22133355                22708822230
Dro. All              : 0                       0
For. All              : 22133355                22708822230

Egress Policer:  7  Group: QGEgr1  Instance-Id: 1
Stats mode: minimal
Off. All              : 22133355                22708822230
Dro. All              : 0                       0
For. All              : 22133355                22708822230

Egress Policer:  8  Group: QGEgr1  Instance-Id: 1
Stats mode: minimal
Off. All              : 22133355                22708822230
Dro. All              : 0                       0
For. All              : 22133355                22708822230
```

```
-------------------------------------------------------------------------------
*A:Dut-T#


*A:Dut-T# show port 9/2/4 queue-group egress "QGEgr1" network associations
===============================================================================
Ethernet port 9/2/4 Network Egress queue-group
===============================================================================
Queue-Group   : QGEgr1             Queue-Id : 1
Queue-Group   : QGEgr1             Policer-*: 1
Queue-Group   : QGEgr1             Policer-*: 2
Queue-Group   : QGEgr1             Policer-*: 3
Queue-Group   : QGEgr1             Policer-*: 4
Queue-Group   : QGEgr1             Policer-*: 5
Queue-Group   : QGEgr1             Policer-*: 6
Queue-Group   : QGEgr1             Policer-*: 7
Queue-Group   : QGEgr1             Policer-*: 8
..
*A:Dut-T#

*A:Dut-T# show qos queue-group "QGIng1" ingress association
===============================================================================
QoS Queue-Group Ingress
===============================================================================
-------------------------------------------------------------------------------
QoS Queue Group
-------------------------------------------------------------------------------
Group-Name     : QGIng1
Description    : Description for Ingress queue-group QGIng1

===============================================================================
Queue Group Ports
===============================================================================
Port                Sched Pol           Acctg Pol Stats     Description
-------------------------------------------------------------------------------
No Matching Entries

===============================================================================
Queue Group Sap FC Maps
===============================================================================
Sap Policy     FC Name             Queue (id type)
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
Queue Group FP Maps
===============================================================================
Card Num       Fp Num              Instance            Type
-------------------------------------------------------------------------------
4              1                   1                   Network
4              1                   2                   Network
4              1                   3                   Network
4              1                   4                   Network
4              1                   5                   Network
4              1                   6                   Network
4              1                   7                   Network
4              1                   8                   Network
4              1                   9                   Network
4              1                   10                  Network
```

```
clear card 4 fp 1 ingress mode network queue-group "QGIng1" instance 1 statistics


*A:Dut-T# monitor card 9 fp 1 ingress queue-group "QGIng1" network instance 1 policer 1
===============================================================================
Monitor Card: 9 Ingress Network Queue-Group: QGIng1 Statistics
===============================================================================
-------------------------------------------------------------------------------
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  1  Grp: QGIng1 (Stats mode: minimal)
Off. All              : 98080861                97685211156
Dro. All              : 12856083                12801694068
For. All              : 85224778                84883517088

-------------------------------------------------------------------------------
At time t = 11 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  1  Grp: QGIng1 (Stats mode: minimal)
Off. All              : 16190                   16125240
Dro. All              : 16010                   15945960
For. All              : 180                     179280

^C
*A:Dut-T#



*A:Dut-T>config>qos>sap-ingress# show card 3 fp 1 ingress queue-group "QGIng3" instance 1
mode access detail
===============================================================================
Card:3  Acc.QGrp: QGIng3  Instance: 1
===============================================================================
Group Name    : QGIng3
Description   : (Not Specified)
Pol Ctl Pol   : None                    Acct Pol      : None
Collect Stats : disabled
-------------------------------------------------------------------------------
Queues
-------------------------------------------------------------------------------
No queues found
-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  1  Grp: QGIng3 (Stats mode: offered-profile-capped-cir)
Off. InProf           :         0                       0
Off. OutProf          :         0                       0
Off. Uncolor          :         22159073                1506816964
Dro. InProf           :         0                       0
Dro. OutProf          :         0                       0
For. InProf           :         215642                  14663656
For. OutProf          :         21943431                1492153308
```

```
Ing. Policer:  2  Grp: QGIng3 (Stats mode: offered-profile-capped-cir)
Off. InProf           :         0                       0
Off. OutProf          :         0                       0
Off. Uncolor          :         274898620               18693106160
Dro. InProf           :         0                       0
Dro. OutProf          :         0                       0
For. InProf           :         1640582                 111559576
For. OutProf          :         273258038               18581546584

Ing. Policer:  3  Grp: QGIng3 (Stats mode: offered-profile-capped-cir)
Off. InProf           :         0                       0
Off. OutProf          :         0                       0
Off. Uncolor          :         19318482                1313656776
Dro. InProf           :         0                       0
Dro. OutProf          :         0                       0
For. InProf           :         188072                  12788896
For. OutProf          :         19130410                1300867880

Ing. Policer:  4  Grp: QGIng3 (Stats mode: offered-profile-capped-cir)
Off. InProf           :         0                       0
Off. OutProf          :         0                       0
Off. Uncolor          :         24634863                1675170684
Dro. InProf           :         0                       0
Dro. OutProf          :         0                       0
For. InProf           :         240244                  16336592
For. OutProf          :         24394619                1658834092
===============================================================================
*A:Dut-T>config>qos>sap-ingress#

*A:Dut-A# show card 9 fp 1 ingress mode access queue-group "Ingress_QG_1" instance 2838
statistics

===============================================================================
Card:9  Acc.QGrp: Ingress_QG_1  Instance: 2838
===============================================================================
Group Name    : Ingress_QG_1
Description   : (Not Specified)
Pol Ctl Pol   : None                    Acct Pol      : None
Collect Stats : disabled
-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              :         53982387                6909745536
Dro. All              :         50861158                6510228224
For. All              :         3121229                 399517312
===============================================================================

*A:Dut-A# show card 9 fp 1 ingress mode access queue-group "Ingress_QG_1" instance 2838

===============================================================================
Card:9  Acc.QGrp: Ingress_QG_1  Instance: 2838
===============================================================================
Group Name    : Ingress_QG_1
Description   : (Not Specified)
Pol Ctl Pol   : None                    Acct Pol      : None
```

```
Collect Stats : disabled
===============================================================================
*A:Dut-A# show card 9 fp 1 ingress mode access queue-group "Ingress_QG_1" instance 2838
detail

===============================================================================
Card:9  Acc.QGrp: Ingress_QG_1  Instance: 2838
===============================================================================
Group Name    : Ingress_QG_1
Description   : (Not Specified)
Pol Ctl Pol   : None                    Acct Pol      : None
Collect Stats : disabled
-------------------------------------------------------------------------------
Queues
-------------------------------------------------------------------------------
No queues found
-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              :         65347348                8364460544
Dro. All              :         61569092                7880843776
For. All              :         3778256                 483616768
===============================================================================
*A:Dut-A# show card 9 fp 1 ingress mode access queue-group "Ingress_QG_1" instance 2838
statistics

===============================================================================
Card:9  Acc.QGrp: Ingress_QG_1  Instance: 2838
===============================================================================
Group Name    : Ingress_QG_1
Description   : (Not Specified)
Pol Ctl Pol   : None                    Acct Pol      : None
Collect Stats : disabled
-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              :         65347348                8364460544
Dro. All              :         61569092                7880843776
For. All              :         3778256                 483616768
===============================================================================

*A:Dut-A# monitor card 9 fp 1 ingress access queue-group "Ingress_QG_1" instance 2838
policer 2

===============================================================================
Monitor Card: 9 Ingress Access Queue-Group: Ingress_QG_1 Statistics
===============================================================================
-------------------------------------------------------------------------------
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
```

```
Off. All              : 133088161               17035284608
Dro. All              : 125393700               16050393600
For. All              : 7694461                 984891008

-------------------------------------------------------------------------------
At time t = 11 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9306452                 1191225856
Dro. All              : 8768431                 1122359168
For. All              : 538021                  68866688

-------------------------------------------------------------------------------
At time t = 22 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9290787                 1189220736
Dro. All              : 8754737                 1120606336
For. All              : 536050                  68614400

-------------------------------------------------------------------------------
At time t = 33 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9291993                 1189375104
Dro. All              : 8753745                 1120479360
For. All              : 538248                  68895744

-------------------------------------------------------------------------------
At time t = 44 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9289980                 1189117440
Dro. All              : 8752910                 1120372480
For. All              : 537070                  68744960

-------------------------------------------------------------------------------
At time t = 55 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9291543                 1189317504
Dro. All              : 8754385                 1120561280
For. All              : 537158                  68756224

-------------------------------------------------------------------------------
At time t = 66 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets
```

```
Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9290688                 1189208064
Dro. All              : 8753578                 1120457984
For. All              : 537110                  68750080

-------------------------------------------------------------------------------
At time t = 77 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9290745                 1189215360
Dro. All              : 8753631                 1120464768
For. All              : 537114                  68750592

-------------------------------------------------------------------------------
At time t = 88 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9290723                 1189212544
Dro. All              : 8753612                 1120462336
For. All              : 537111                  68750208

-------------------------------------------------------------------------------
At time t = 99 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9290589                 1189195392
Dro. All              : 8753485                 1120446080
For. All              : 537104                  68749312

-------------------------------------------------------------------------------
At time t = 110 sec (Mode: Delta)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ing. Policer:  2  Grp: Ingress_QG_1 (Stats mode: minimal)
Off. All              : 9290735                 1189214080
Dro. All              : 8753622                 1120463616
For. All              : 537113                  68750464

===============================================================================


*A:Dut-A# clear card 9 fp 1 mode access ingress queue-group "Ingress_QG_1" instance 2838
statistics
```

sap-egress

Syntax **sap-egress** [*policy-id*] [**association** | **match-criteria** | **hsmda** | **detail**]

Context show>qos

Description This command displays SAP egress QoS policy information. Queue group information is displayed in the FC section.

Parameters *policy-id* — The SAP egress policy ID that uniquely identifies the policy.

association — Displays the entities associated with the specified policy ID.

match-criteria — Displays match criteria when this keyword is specified.

hsmda — Displays HSMDA properties.

detail — Displays detailed information about the specified SAP egress policy.

Sample Output

```
*A:Dut-T>config>port# show qos sap-egress 10 detail
===============================================================================
QoS Sap Egress
===============================================================================
-------------------------------------------------------------------------------
Sap Egress Policy (10)
-------------------------------------------------------------------------------
Policy-id      : 10                           Scope        : Template
Description    : (Not Specified)

-------------------------------------------------------------------------------
Queue CIR Admin PIR Admin CBS     HiPrio PIR Lvl/Wt    Parent         AvgOvrhd
      CIR Rule  PIR Rule  MBS            CIR Lvl/Wt
      Named-Buffer Pool
-------------------------------------------------------------------------------
1     0         max       def     def    1/1           None           0.00
      closest   closest   def            0/1
      (not-assigned)


-------------------------------------------------------------------------------
FC Name   Queue QGroup   Dot1p Exp/Default        DE-Mark DSCP/Prec Marking
-------------------------------------------------------------------------------
be        1    QG_egres* Default                  None    default
af        2    QG_egres* Default                  None    default
ef        3    QG_egres* Default                  None    default
nc        4    QG_egres* Default                  None    default

-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 1 (VPLS)                     Customer-Id  : 1
 - SAP : 9/2/1
 - SAP : 9/2/2

-------------------------------------------------------------------------------
```

```
Mirror SAPs
-------------------------------------------------------------------------------
No Mirror SAPs Found.


-------------------------------------------------------------------------------
HSMDA CIR Admin PIR Admin Packet  Slope Policy
Queue CIR Rule  PIR Rule  Offset
-------------------------------------------------------------------------------
1     0         max       add 0   default
      closest   closest
2     0         max       add 0   default
      closest   closest
3     0         max       add 0   default
      closest   closest
4     0         max       add 0   default
      closest   closest
5     0         max       add 0   default
      closest   closest
6     0         max       add 0   default
      closest   closest
7     0         max       add 0   default
      closest   closest
8     0         max       add 0   default
      closest   closest

-------------------------------------------------------------------------------
FC              HSMDA Queue-id               HSMDA Dot1p Profiling
-------------------------------------------------------------------------------
af              def                          disabled
be              def                          disabled
ef              def                          disabled
nc              def                          disabled

-------------------------------------------------------------------------------
DSCP            Cntr Id   Profile   fc
-------------------------------------------------------------------------------
No DSCP-Map Entries Found.


-------------------------------------------------------------------------------
Prec Value      Cntr Id   Profile   fc
-------------------------------------------------------------------------------
No Prec-Map Entries Found.

-------------------------------------------------------------------------------
Match Criteria
-------------------------------------------------------------------------------
No Matching Criteria.


-------------------------------------------------------------------------------
HSMDA Associations
-------------------------------------------------------------------------------
No Associations Found.

===============================================================================
*A:Dut-T>config>port#
```

sap-ingress

Syntax **sap-ingress** [*policy-id*] [**association** | **match-criteria** | **hsmda** | **detail**]

Context show>qos

Description This command displays SAP ingress QoS policy information. Queue group information is displayed in the FC section.

Parameters *policy-id* — The SAP egress policy ID that uniquely identifies the policy.

association — Displays the entities associated with the specified policy ID.

match-criteria — Displays match criteria when this keyword is specified.

hsmda — Displays HSMDA properties.

detail — Displays detailed information about the specified SAP egress policy.

Sample Output

```
*A:Dut-T>config>port# show qos sap-ingress 10 detail
===============================================================================
QoS Sap Ingress
===============================================================================
-------------------------------------------------------------------------------
Sap Ingress Policy (10)
-------------------------------------------------------------------------------
Policy-id      : 10                           Scope        : Template
Default FC     : be                           Priority     : Low
Criteria-type  : None
Description    : (Not Specified)

-------------------------------------------------------------------------------
Queue Mode    CIR Admin PIR Admin CBS     HiPrio  PIR Lvl/Wt    Parent
              CIR Rule  PIR Rule  MBS             CIR Lvl/Wt
              Named-Buffer Pool
-------------------------------------------------------------------------------
1     Prio      0         max     def    def           1/1            None
                closest   closest def                  0/1
                (not-assigned)
11    Prio      0         max     def    def           1/1            None
                closest   closest def                  0/1
                (not-assigned)

-------------------------------------------------------------------------------
FC   UCastQ/QGrp   MCastQ/QGrp    BCastQ/QGrp   UnknownQ/QGrp
-------------------------------------------------------------------------------
be              1/QG_ingress_1 def/           def/           def/
af              2/QG_ingress_1 def/           def/           def/
ef              3/QG_ingress_1 def/           def/           def/
nc              4/QG_ingress_1 def/           def/           def/

-------------------------------------------------------------------------------
FC           DE-1-out-profile Profile        In-Remark      Out-Remark
-------------------------------------------------------------------------------
af              No               None           None           None
```

```
be              No              None          None          None
ef              No              None          None          None
nc              No              None          None          None

-------------------------------------------------------------------------------
Dot1p           FC                                Priority            Cntr Id
-------------------------------------------------------------------------------
0               be                                Default             Default
2               af                                Default             Default
4               ef                                Default             Default
6               nc                                Default             Default

-------------------------------------------------------------------------------
DSCP            FC                                Priority            Cntr Id
-------------------------------------------------------------------------------
No DSCP-Map Entries Found.


-------------------------------------------------------------------------------
Prec Value      FC                                Priority            Cntr Id
-------------------------------------------------------------------------------
No Prec-Map Entries Found.

-------------------------------------------------------------------------------
Match Criteria
-------------------------------------------------------------------------------
No Matching Criteria.

-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 1 (VPLS)                    Customer-Id  : 1
 - SAP : 9/2/1
 - SAP : 9/2/2

-------------------------------------------------------------------------------
HSMDA CIR Admin PIR Admin Packet  Slope Policy
Queue CIR Rule  PIR Rule  Offset
-------------------------------------------------------------------------------
1     0         max       add 0   default
      closest   closest
2     0         max       add 0   default
      closest   closest
3     0         max       add 0   default
      closest   closest
4     0         max       add 0   default
      closest   closest
5     0         max       add 0   default
      closest   closest
6     0         max       add 0   default
      closest   closest
7     0         max       add 0   default
      closest   closest
8     0         max       add 0   default
      closest   closest

-------------------------------------------------------------------------------
FC              HSMDA UCastQ   HSMDA MCastQ   HSMDA BCastQ
-------------------------------------------------------------------------------
```

```
af              def            def            def
be              def            def            def
ef              def            def            def
nc              def            def            def
-------------------------------------------------------------------------------
HSMDA Associations
-------------------------------------------------------------------------------
No Associations Found.
===============================================================================
*A:Dut-T>config>port#
```

pools

Syntax **pools** *mda-id*[*/port*] [*access-app* [*pool-name* | **service** *service-id* | **queue-group** *queue-group-name*]]

pools *mda-id*[*/port*] [**network-app** [*pool-name* | **queue-group** *queue-group-name*]]

pools *mda-id*[*/port*] [**direction** [*pool-name*|**service** *service-id* | **queue-group** *queue-group-name*]]

Context show

Description This command displays queue group pool information.

Parameters *mda-id*[*/port*] — Displays the pool information of the specified MDA.

access-app pool-name — Displays the pool information of the specified QoS policy.

Values access-ingress, access-egress

service *service-id* **—** Displays pool information for the specified service.

Values 1 — 2147483647

queue-group *queue-group-name* **—** Display information for the specified queue group.

direction — Specifies to display information for the ingress or egress direction.

Values ingress, egress

Sample Output

```
*A:Dut-T>config>port# show pools 9/2/1 access-egress  queue-group QG_egress_1

===============================================================================
Pool Information
===============================================================================
Port                 : 9/2/1
Application          : Acc-Egr           Pool Name          : default
Resv CBS             : Sum
-------------------------------------------------------------------------------
Queue-Groups
-------------------------------------------------------------------------------
QG_egress_1
-------------------------------------------------------------------------------
Utilization                  State       Start-Avg    Max-Avg    Max-Prob
```

```
-------------------------------------------------------------------------------
High-Slope                    Down              70%        90%         80%
Low-Slope                     Down              50%        75%         80%

Time Avg Factor      : 7
Pool Total           : 6336 KB
Pool Shared          : 4416 KB           Pool Resv          : 1920 KB

Pool Total In Use    : 0 KB
Pool Shared In Use   : 0 KB              Pool Resv In Use   : 0 KB
WA Shared In Use     : 0 KB

Hi-Slope Drop Prob   : 0                 Lo-Slope Drop Prob : 0
-------------------------------------------------------------------------------
Name                          FC-Maps       MBS       HP-Only A.PIR    A.CIR
                                            CBS       Depth   O.PIR    O.CIR
-------------------------------------------------------------------------------
QGrp->QG_egress_1(9/2/1)->1
                              n/a           102       9       1000000  0
                                            0         0       Max      0
QGrp->QG_egress_1(9/2/1)->2
                              n/a           102       9       1000000  0
                                            0         0       Max      0
QGrp->QG_egress_1(9/2/1)->3
                              n/a           102       9       1000000  0
                                            0         0       Max      0
QGrp->QG_egress_1(9/2/1)->4
                              n/a           102       9       1000000  0
                                            0         0       Max      0
===============================================================================
*A:Dut-T>config>port#


*A:Dut-T>config>port# show pools 9/2/1 access-ingress queue-group QG_ingress_1
===============================================================================
Pool Information
===============================================================================
Port                 : 9/2/1
Application          : Acc-Ing           Pool Name          : default
Resv CBS             : Sum
-------------------------------------------------------------------------------
Queue-Groups
-------------------------------------------------------------------------------
QG_ingress_1
-------------------------------------------------------------------------------
Utilization                   State       Start-Avg    Max-Avg    Max-Prob
-------------------------------------------------------------------------------
High-Slope                    Down              70%        90%         80%
Low-Slope                     Down              50%        75%         80%

Time Avg Factor      : 7
Pool Total           : 168960 KB
Pool Shared          : 116736 KB         Pool Resv          : 52224 KB

Pool Total In Use    : 0 KB
Pool Shared In Use   : 0 KB              Pool Resv In Use   : 0 KB
WA Shared In Use     : 0 KB

Hi-Slope Drop Prob   : 0                 Lo-Slope Drop Prob : 0
```

```
-------------------------------------------------------------------------------
Name                            FC-Maps       MBS       HP-Only A.PIR    A.CIR
                                              CBS       Depth   O.PIR    O.CIR
-------------------------------------------------------------------------------
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a           102       9       1000000  0
                                              0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
```

```
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->1
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->2
```

```
QGrp->QG_ingress_1(9/2/1)->2
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->2
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->3
                                n/a             102       9       1000000  0
```

```
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->3
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->3
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
                              n/a       102       9       1000000   0
                                        0         0       Max       0
QGrp->QG_ingress_1(9/2/1)->4
```

```
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->4
                                n/a             102       9       1000000  0
                                                0         0       Max      0
QGrp->QG_ingress_1(9/2/1)->4
                                n/a             102       9       1000000  0
                                                0         0       Max      0
===============================================================================
*A:Dut-T>config>port#
```

port

Syntax **port** *port-id* **queue-group** [**ingress** | **egress**] [*queue-group-name*][{**statistics** | **associations**}]

Context show>port

Description This command displays physical port information for the port's queue group.

Parameters *port-id* — Specifies the port ID to display information abou the port's queue group.

queue-group ingress — Specifies whether the queue group name is an ingress policy.

queue-group egress — Specifies whether the queue group name is an egress policy.

queue-group-name — Specifies the name of an existing queue group template up to 32 characters in length.

statistics — Displays statistical information for the queue group.

associations — Displays the entities associated with the specified queue group name.

Sample Output

```
*A:Dut-T>config>port# show port 9/2/1 queue-group ingress
===============================================================================
Ethernet port 9/2/1 Access Ingress queue-group
===============================================================================
Group Name    : QG_ingress_1
Description   : (Not Specified)
Sched Policy  : None               Acct Pol : None
Collect Stats : disabled

Queues
-------------------------------------------------------------------------------
Ing. QGroup   : QG_ingress_1       Queue-Id : 1 (Unicast) (Priority)
Description   : Ingress queue-group
Admin PIR     : max*               Admin CIR: 0*
PIR Rule      : closest*           CIR Rule : closest*
CBS           : def*               MBS      : 100*
Hi Prio       : def*

Ing. QGroup   : QG_ingress_1       Queue-Id : 2 (Unicast) (Priority)
Description   : Ingress queue-group
Admin PIR     : 800000             Admin CIR: 20000
PIR Rule      : closest*           CIR Rule : closest*
CBS           : def*               MBS      : 100*
Hi Prio       : def*

Ing. QGroup   : QG_ingress_1       Queue-Id : 3 (Unicast) (Priority)
Description   : Ingress queue-group
Admin PIR     : max*               Admin CIR: 0*
PIR Rule      : closest*           CIR Rule : closest*
CBS           : def*               MBS      : 100*
Hi Prio       : def*

Ing. QGroup   : QG_ingress_1       Queue-Id : 4 (Unicast) (Priority)
```

```
Description   : Ingress queue-group
Admin PIR     : max*               Admin CIR: 0*
PIR Rule      : closest*           CIR Rule : closest*
CBS           : def*               MBS      : 100*
Hi Prio       : def*

* means the value is inherited
===============================================================================
*A:Dut-T>config>port#


*A:Dut-T>config>port# show port 9/2/2 queue-group egress
===============================================================================
Ethernet port 9/2/2 Access Egress queue-group
===============================================================================
Group Name    : QG_egress_1
Description   : (Not Specified)
Sched Policy  : None               Acct Pol : None
Collect Stats : disabled

Queues
-------------------------------------------------------------------------------
Egr. QGroup   : QG_egress_1        Queue-Id : 1
Description   : Egress queue group
Admin PIR     : max*               Admin CIR: 0*
PIR Rule      : closest*           CIR Rule : closest*
CBS           : def*               MBS      : 100*
Hi Prio       : def*

Egr. QGroup   : QG_egress_1        Queue-Id : 2
Description   : Egress queue group
Admin PIR     : max*               Admin CIR: 0*
PIR Rule      : closest*           CIR Rule : closest*
CBS           : def*               MBS      : 100*
Hi Prio       : def*

Egr. QGroup   : QG_egress_1        Queue-Id : 3
Description   : Egress queue group
Admin PIR     : 1500000            Admin CIR: 2000
PIR Rule      : closest*           CIR Rule : closest*
CBS           : def*               MBS      : 100*
Hi Prio       : def*

Egr. QGroup   : QG_egress_1        Queue-Id : 4
Description   : Egress queue group
Admin PIR     : max*               Admin CIR: 0*
PIR Rule      : closest*           CIR Rule : closest*
CBS           : def*               MBS      : 100*
Hi Prio       : def*

* means the value is inherited
===============================================================================
*A:Dut-T>config>port#

*A:Dut-T>config>port# show port 9/2/2 egress queue-group QG_egress_1 statistics
-------------------------------------------------------------------------------
Ethernet port 9/2/2 Access Egress queue-group
-------------------------------------------------------------------------------
                        Packets                 Octets
```

```
Egress Queue: 1 Group: QG_egress_1
For. InProf           : 0                     0
For. OutProf          : 228091788             14959815064
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0

Egress Queue: 2 Group: QG_egress_1
For. InProf           : 0                     0
For. OutProf          : 40661626              2764990568
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0

Egress Queue: 3 Group: QG_egress_1
For. InProf           : 0                     0
For. OutProf          : 40661628              2764990704
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0

Egress Queue: 4 Group: QG_egress_1
For. InProf           : 0                     0
For. OutProf          : 40661629              2764990772
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0
-------------------------------------------------------------------------------
*A:Dut-T>config>port#
```

Monitor Commands

card

Syntax **card** *slot-number* **fp** *fp-number* **ingress** {**access** | **network**} **queue-group** *queue-group-name* **instance** *instance-id* [**interval** *seconds*] [**repeat** *repeat*] **policer** *policer-id* [**absolute** | **percent-rate** | *reference-rate*]

Context monitor

Description This command monitors policer statistics in an ingress FP queue group.

Parameters **card** *slot-number* **—** Specifies the slot number associated with the queue group, expressed as an integer.

Values 1 — 20

fp *fp-number* **—** Specifies the FP number associated with the queue group, expressed as an integer .

Values 1 — 2

ingress — Displays policer statistics applied on the ingress FP.

access — Displays policer statistics on the FP access.

network — Displays policer statistics on the FP network.

queue-group *queue-group-name* **—** Specifies the name of the queue group up to 32 characters in length

instance *instance-id* **—** Specifies the identification of a specific instance of the queue-group.

Values 1— 65535

interval *seconds* **—** Configures the interval for each display in seconds.

Default 11 seconds

Values 11 — 60

repeat *repeat* **—** Configures how many times the command is repeated.

Default 10

Values 1 — 999

policer *policer-id* **—** The specified policer-id must exist within the queue-group template applied to the ingress context of the forwarding plane.

Values 1—32

absolute — When the absolute keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

percent-rate — When the percent-rate keyword is specified, the rate-per-second for each statistic is displayed based on the reference rate of 10G.

reference-rate — When a reference-rate value is specified, the rate-per-second for each statistic is displayed as a percentage based on the reference rate specified.

Values 100M, 1G, 10G, 40G, 100G, 400G

card

Syntax **card** *slot-number* **fp** *fp-number* **queue-group** *queue-group-name* **instance** *instance-id* [**ingress**] [**access** | **networks**] [**interval** *seconds*] [**repeat** *repeat***]** [**absolute** | **percent-rate |** *reference-rate*] [**arbiter** *root* **|** *name*]

Context monitor>qos>arbiter-stats

Description This command monitors arbiter statistics in an ingress FP queue group.

Parameters **card** *slot-number* **—** Specifies the slot number associated with the queue group, expressed as an integer.

Values 1 — 20

fp *fp-number* **—** Specifies the FP number associated with the queue group, expressed as an integer .

Values 1 — 2

ingress — Displays policer statistics applied on the ingress FP.

access — Displays policer statistics on the FP access.

network — Displays policer statistics on the FP network.

queue-group *queue-group-name* **—** Specifies the name of the queue group up to 32 characters in length

instance *instance-id* **—** Specifies the identification of a specific instance of the queue-group.

Values 1— 65535

interval *seconds* **—** Configures the interval for each display in seconds.

Default 11 seconds

Values 11 — 60

repeat *repeat* **—** Configures how many times the command is repeated.

Default 10

Values 1 — 999

absolute — When the absolute keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

percent-rate — When the percent-rate keyword is specified, the rate-per-second for each statistic is displayed based on the reference rate of 10G.

reference-rate — When a reference-rate value is specified, the rate-per-second for each statistic is displayed as a percentage based on the reference rate specified.

Values 100M, 1G, 10G, 40G, 100G, 400G

arbiter name — — Specify the name of the policer control policy arbiter.

Values An existing arbiter-name in the form of a string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

root — Specifies the root arbiter.

port

Syntax **port** *port-id* **egress** *network* **queue-group** *queue-group-name* **instance** *instance-id* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] [**arbiter** *root* | *name*]

Context monitor>qos>arbiter-stats

Description This command monitors arbiter statistics in an egress port queue group.

Parameters **port** *port-id* **—** Specifies the port ID.

Values slot/mda/port

egress *network* **—** Specifies statistics are for an egress network queue group.

queue-group *queue-group-name* **—** Specifies the name of the queue group up to 32 characters in length.

instance *instance-id* **—** Specifies the identification of a specific instance of the queue-group.

Values 1— 65535

interval *seconds* **—** Configures the interval for each display in seconds.

Default 11 seconds

Values 11 — 60

repeat *repeat* **—** Configures how many times the command is repeated.

Default 10

Values 1 — 999

absolute — When the **absolute** keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

rate — When the **rate** keyword is specified, the rate-per-second for each statistic is displayed.

arbiter name — Specify the name of the policer control policy arbiter.

Values An existing arbiter-name in the form of a string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

root — Specify the root arbiter.

port

Syntax **port** *port-id* **queue-group** *queue-group-name* [**ingress** | **egress**] [**interval** *seconds*] [**repeat** *repeat*] [**absolute | rate**] [**access** | **network**] **[instance** *instance-id*]

Context monitor>qos>scheduler-stats

Description This command enables port traffic monitoring.

Parameters **port** *port-id* **—** Specifies the port ID.

Values slot/mda/port

queue-group *queue-group-name* **—** Specifies the name of the queue group up to 32 characters in length.

instance *instance-id* **—** Specifies the identification of a specific instance of the queue-group.

Values 1— 65535

ingress — Specifies statistics are for an ingress queue group.

egress — Specifies statistics are for an egress queue group.

interval *seconds* **—** Configures the interval for each display in seconds.

Default 11 seconds

Values 11 — 60

repeat *repeat* **—** Configures how many times the command is repeated.

Default 10

Values 1 — 999

absolute — When the **absolute** keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

rate — When the **rate** keyword is specified, the rate-per-second for each statistic is displayed.

access — Displays scheduler statistics applied on an access port.

network — Displays scheduler statistics applied on a network port.

QoS Queue Group Template Command Reference

Command Hierarchies

- Configuring Egress Queue Group Templates on page 495
- Configuring Ingress Queue Group Templates on page 497
- Show Commands on page 498
- Monitor Commands on page 498

Configuring Egress Queue Group Templates

```
config
    — qos
        — queue-group-templates
            — egress
                — queue-group queue-group-name [create]
                — no queue-group queue-group-name
                    — description description-string
                    — no description
                    — policer policer-id [create]
                    — no policer policer-id
                        — adaptation-rule pir {max | min | closest}] [cir {max | min | closest}]
                        — no adaptation-rule
                        — adv-config-policy policy-name
                        — no adv-config-policy
                        — description description-string
                        — no description
                        — cbs {size [bytes | kilobytes] | default}
                        — no cbs
                        — high-prio-only percent-of-mbs
                        — no high-prio-only
                        — mbs {size [bytes | kilobytes] | default}
                        — no mbs
                        — packet-byte-offset {add bytes | subtract bytes}
                        — no packet-byte-offset
                        — parent {root | arbiter-name} [level level] [weight weight-within-level]
                        — no parent
                        — profile-cappedrate {max | kilobits-per-second} [cir {max | kilobits-per-second}]
                        — no rate
                        — stat-mode {no-stats | minimal | offered-profile-no-cir | offered-priority-no-cir | offered-limited-profile-cir | offered-profile-cir | offered-priority-cir | offered-total-cir | offered-profile-capped-cir | offered-limited-capped-cir}
                        — no stat-mode
                    — queue queue-id [queue-type] [create]
                    — no queue queue-id
```

- **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
- **no adaptation-rule**
- **adv-config-policy** *adv-config-policy-name*
- **no adv-config-policy**
- **burst-limit**
- **no burst-limit**
- **burst-limit** *size-in-kbytes*
- **no burst-limit**
- **high-prio-only** *percent*
- **no high-prio-only**
- **mbs** *size* [**bytes** | **kilobytes**]
- **no mbs**
- **parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
- **no parent**
- **percent-rate** *per-percent* [**cir** *cir-percent*]
- **no percent-rate**
- **port-parent** [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
- **no port-parent**
- **rate** *pir-rate* [**cir** *cir-rate*]
- **no rate**
  - **packet-byte-offset** [**policy** *slope-policy-name*]
  - **no packet-byte-offset**
- **xp-specific**
  - **packet-byte-offset {add** *bytes* | **subtract** *bytes***}**

Configuring Ingress Queue Group Templates

config
- **qos**
  - **queue-group-templates**
    - **ingress**
      - **description** *description-string*
      - **no description**
      - **queue-group** *queue-group-name* [**create**]
      - **no queue-group** *queue-group-name*
        - **policer** *policer-id* [**create**]
        - [**no**] **policer**
          - **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
          - **no adaptation-rule**
          - **cbs** {*size* [**bytes** | **kilobytes**] | **default**}
          - **no cbs**
          - **description** *description string*
          - **no description**
          - **high-prio-only** *percent-of-mbs*
          - **no high-prio-only**
          - **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
          - **no mbs**
          - **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}
          - **no packet-byte-offset**
          - **parent** {**root** | *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]
          - **no parent**
          - **rate** {**max** | **kilobits-per-second**} [**cir** {**max** | **kilobits-per-second**}]
          - **no rate**
          - **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-profile-cir** | **offered-total-cir**}
          - **no stat-mode**
        - **queue** *queue-id* [**multipoint**] [*queue-type*] [*queue-mode*] [**create**]
        - **no queue** *queue-id*
          - **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
          - **no adaptation-rule**
          - **burst-limit**
          - **no burst-limit**
          - **burst-limit** *size-in-kbytes*
          - **no burst-limit**
          - **high-prio-only** *percent*
          - **no high-prio-only**
          - **mbs** *size-in-kbytes*
          - **no mbs**
          - **parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
          - **no parent**
          - **rate** *pir-rate* [**cir** *cir-rate*]
          - **rate** *pir-rate* **police**
          - **no rate**

Show Commands

show
- **qos**
  - **queue-group** [*queue-group-name*] [**ingress** | **egress**] [**association** | **detail**]
  - **queue-group summary**
  - **sap-egress** [*policy-id*] [**association** | **match-criteria**| **hsmda** | **detail**]
  - **sap-ingress** [*policy-id*] [**association** | **match-criteria**| **hsmda** | **detail**]

show
- **pools** *mda-id*[/*port*] [*access-app* [*pool-name* | **service** *service-id* | **queue-group** *queue-group-name*]]
- **pools** *mda-id*[/*port*] [**network-app** [*pool-name* | **queue-group** *queue-group-name*]]
- **pools** *mda-id*[/*port*] [**direction** [*pool-name* | **service** *service-id* | **queue-group** *queue-group-name*]]
- **port** *port-id* **queue-group** [**ingress** | **egress**] [*queue-group-name*][{**statistics** | **associations**}]

Monitor Commands

monitor
- **card slot-number fp** *fp-number* **ingress** {**access** | **network**} **queue-group** *queue-group-name* **instance** *instance-id* **[interval** *seconds*]**[repeat** *repeat*] **policer** *policer-id* **[absolute | percent-rate |** *reference-rate*] [**arbiter** root | name]
- **qos**
  - **arbiter-stats**
    - **card** *slot-number* **fp** *fp-number* **queue-group** *queue-group-name* **instance** *instance-id* [**ingress**] [**access** | **networks**] [**interval** *seconds*][**repeat** *repeat*] [**absolute | rate**] [**arbiter** *root* | *name*]
    - **port** *port-id* **egress** *network* **queue-group** *queue-group-name* **instance** *instance-id* [**interval** *seconds*] [**repeat** *repeat*] [**absolute | rate**] [**arbiter** *root* | *name*]
  - **scheduler-stats**
    - **port** *port-id* **queue-group** *queue-group-name* [**ingress** | **egress**] [**interval** *seconds*] [**repeat** *repeat*] [**absolute | rate**] [**access** | **network**] **[instance** *instance-id*]

QoS Scheduler Policies

In This Section

This section provides information to configure QoS scheduler and port scheduler policies using the command line interface.

Topics in this section include:

- Overview on page 500
- Basic Configurations on page 524
- Service Management Tasks on page 545

Overview

Scheduler Policies

Virtual schedulers are created within the context of a scheduler policy that is used to define the hierarchy and parameters for each scheduler. A scheduler is defined in the context of a tier which is used to place the scheduler within the hierarchy. Three tiers of virtual schedulers are supported. Root schedulers are defined without a parent scheduler meaning it is not subject to obtaining bandwidth from a higher tier scheduler. A scheduler has the option of enforcing a maximum rate of operation for all child queues and schedulers associated with it.

Because a scheduler is designed to arbitrate bandwidth between many inputs, a metric must be assigned to each child queue or scheduler vying for transmit bandwidth. This metric indicates whether the child is to be scheduled in a strict or weighted fashion and the level or weight the child has to other children.

Egress Port-Based Schedulers

In previous releases, HQoS root (top tier) schedulers always assumed that the configured rate was available, regardless of egress port level oversubscription and congestion. This resulted in the possibility that the aggregate bandwidth assigned to queues was not actually available at the port level. When the HQoS algorithm configures queues with more bandwidth than available on an egress port, actual bandwidth distribution to queues on the port will be solely based on the action of the hardware scheduler. This can result in a forwarding rate at each queue that is very different than the desired rate.

The port-based scheduler feature was introduced to allow HQoS bandwidth allocation based on available bandwidth at the egress port level. The port-based scheduler works at the egress line rate of the port to which it is attached. Port-based scheduling bandwidth allocation automatically includes the Inter-Frame Gap (IFG) and preamble for packets forwarded on queues servicing egress Ethernet ports. However, on PoS and SDH based ports, the HDLC encapsulation overhead and other framing overhead per packet is not known by the system. Instead of automatically determining the encapsulation overhead for SDH or SONET queues, the system provides a configurable frame encapsulation efficiency parameter that allows the user to select the average encapsulation efficiency for all packets forwarded out the egress queue.

A special port scheduler policy can be configured to define the virtual scheduling behavior for an egress port. The port scheduler is a software-based state machine managing a bandwidth allocation algorithm that represents the scheduling hierarchy shown in Figure 17 on page 503.

The first tier of the scheduling hierarchy manages the total frame based bandwidth that the port scheduler will allocate to the eight priority levels.

The second tier receives bandwidth from the first tier in two priorities, a “within-cir” loop and an “above-cir” loop. The second tier “within-cir” loop provides bandwidth to the third tier “within-cir” loops, one for each of the eight priority levels. The second tier “above-cir” loop provides bandwidth to the third tier “above-cir” loops for each of the eight priority levels.

The “within-cir” loop for each priority level on the third tier supports an optional rate limiter used to restrict the maximum amount of “within-cir” bandwidth the priority level can receive. A maximum priority level rate limit is also supported that restricts the total amount of bandwidth the level can receive for both “within-cir” and “above-cir”. The amount of bandwidth consumed by each priority level for “within-cir” and “above-cir” is predicated on the rate limits described and the ability for each child queue or scheduler attached to the priority level to use the bandwidth.

The priority 1 “above-cir” scheduling loop has a special two tier strict distribution function. The high priority level 1 “above-cir” distribution is weighted between all queues and schedulers attached to level 1 for “above-cir” bandwidth. The low priority distribution for level 1 “above-cir” is reserved for all orphaned queues and schedulers on the egress port. Orphans are queues and schedulers that are not explicitly or indirectly attached to the port scheduler through normal parenting conventions. By default, all orphans receive bandwidth after all parented queues and schedulers and are allowed to consume whatever bandwidth is remaining. This default behavior for orphans can be overridden on each port scheduler policy by defining explicit orphan port parent association parameters.

Ultimately, any bandwidth allocated by the port scheduler is given to a child queue. The bandwidth allocated to the queue is converted to a value for the queue’s PIR (maximum rate) setting. This way, the hardware schedulers operating at the egress port level will only schedule bandwidth for all queues on the port up to the limits prescribed by the virtual scheduling algorithm.

The following lists the bandwidth allocation sequence for the port virtual scheduler:

1. Priority level 8 offered load up to priority CIR
2. Priority level 7 offered load up to priority CIR
3. Priority level 6 offered load up to priority CIR
4. Priority level 5 offered load up to priority CIR
5. Priority level 4 offered load up to priority CIR
6. Priority level 3 offered load up to priority CIR
7. Priority level 2 offered load up to priority CIR
8. Priority level 1 offered load up to priority CIR
9. Priority level 8 remaining offered load up to remaining priority rate limit
10. Priority level 7 remaining offered load up to remaining priority rate limit

11. Priority level 6 remaining offered load up to remaining priority rate limit
12. Priority level 5 remaining offered load up to remaining priority rate limit
13. Priority level 4 remaining offered load up to remaining priority rate limit
14. Priority level 3 remaining offered load up to remaining priority rate limit
15. Priority level 2 remaining offered load up to remaining priority rate limit
16. Priority level 1 remaining offered load up to remaining priority rate limit
17. Priority level 1 remaining orphan offered load up to remaining priority rate limit (default orphan behavior unless orphan behavior has been overridden in the scheduler policy)

When a queue is inactive or has a limited offered load that is below its fair share (fair share is based on the bandwidth allocation a queue would receive if it was registering adequate activity), its operational PIR must be set to some value to handle what would happen if the queues offered load increased prior to the next iteration of the port virtual scheduling algorithm. If an inactive queues PIR was set to zero (or near zero), the queue would throttle its traffic until the next algorithm iteration. If the operational PIR was set to its configured rate, the result could overrun the expected aggregate rate of the port scheduler.

To accommodate inactive queues, the system calculates a Minimum Information Rate (MIR) for each queue. To calculate each queue's MIR, the system determines what that queue's Fair Information Rate (FIR) would be if that queue had actually been active during the latest iteration of the virtual scheduling algorithm. For example, if three queues are active (1, 2, and 3) and two queues are inactive (4 and 5), the system first calculates the FIR for each active queue. Then it recalculates the FIR for queue 4 assuming queue 4 was active with queues 1, 2, and 3 and uses the result as the queue's MIR. The same is done for queue 5 using queues 1, 2, 3, and 5. The MIR for each inactive queue is used as the operational PIR for each queue.

Service/Subscriber Egress Port Bandwidth Allocation

The port-based egress scheduler can be used to allocate bandwidth to each service or subscriber associated with the port. While egress queues on the service can have a child association with a scheduler policy on the SAP or multi-service site, all queues must vie for bandwidth from an egress port. Two methods are supported to allocate bandwidth to each service or subscriber queue:

1. Service or subscriber queue association with a scheduler on the SAP or multi-service site which is itself associated with a port-level scheduler.
2. Service or subscriber queue association directly with a port-level scheduler.

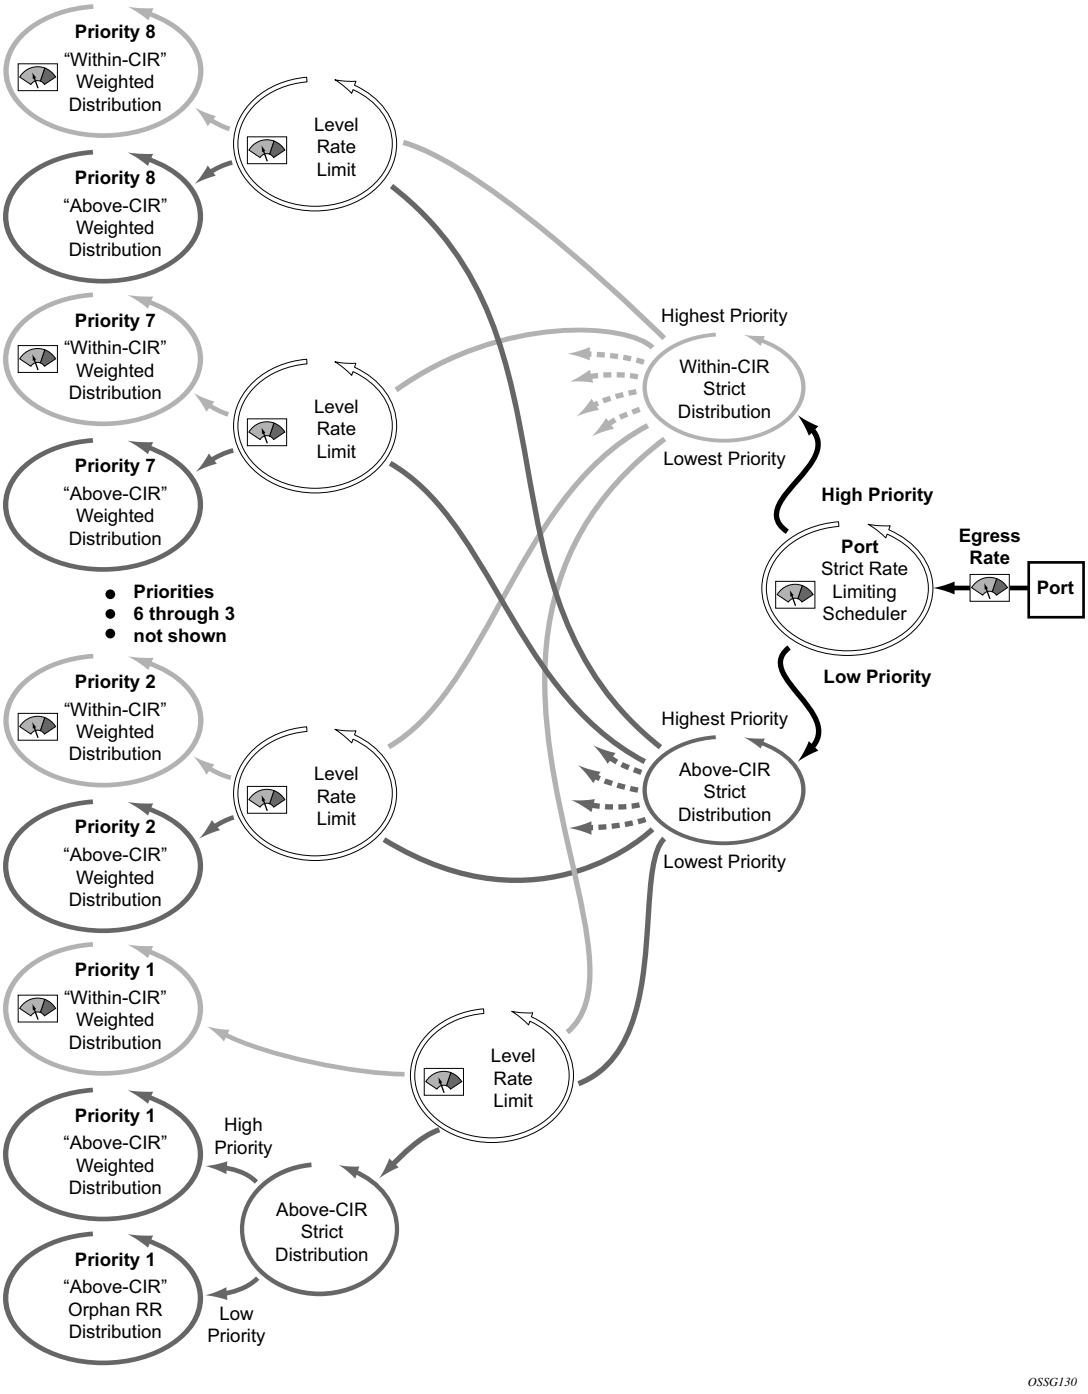

Figure 17: Port Level Virtual Scheduler Bandwidth Allocation Based on Priority and CIR

Service or Subscriber Scheduler Child to Port Scheduler Parent

The service or subscriber scheduler to port scheduler association model allows for multiple services or subscriber to have independent scheduler policy definitions while the independent schedulers receive bandwidth from the scheduler at the port level. By using two scheduler policies, available egress port bandwidth can be allocated fairly or unfairly depending on the desired behavior. Figure 18 graphically demonstrates this model.

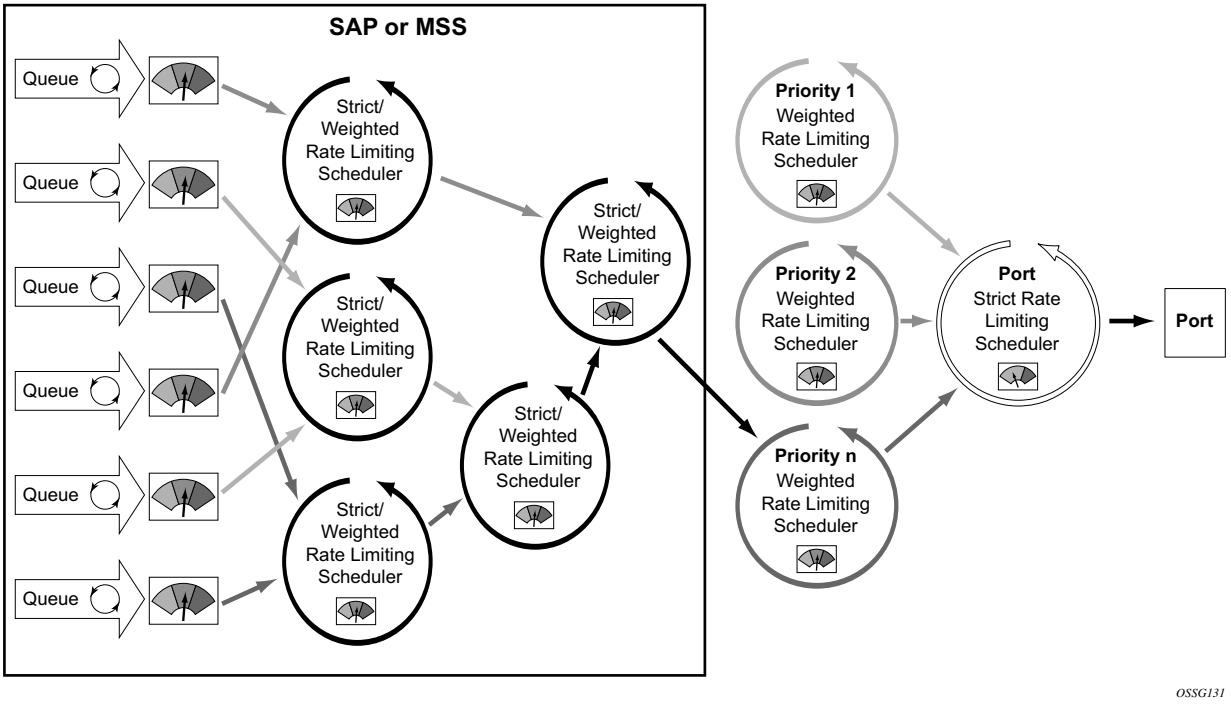

Figure 18: Two Scheduler Policy Model for Access Ports

Once a two scheduler policy model is defined, the bandwidth distribution hierarchy allocates the available port bandwidth to the port schedulers based on priority, weights, and rate limits. The service or subscriber level schedulers and the queues they service become an extension of this hierarchy.

Due to the nature of the two scheduler policy, bandwidth is allocated on a per-service or per subscriber basis as opposed to a per-class basis. A common use of the two policy model is for a carrier-of-carriers mode of business. In essence, the goal of a carrier is to provide segments of bandwidth to providers who purchase that bandwidth as services. While the carrier does not concern itself with the interior services of the provider, it does however care how congestion affects the bandwidth allocation to each provider's service. As an added benefit, the two policy approach provides the carrier with the ability to preferentially allocate bandwidth within a service

or subscriber context through the service or subscriber level policy without affecting the overall bandwidth allocation to each service or subscriber. Figure 19 shows a per-service bandwidth allocation using the two scheduler policy model. While the figure shows services grouped by scheduling priority, it is expected that many service models will place the services in a common port priority and use weights to provide a weighted distribution between the service instances. Higher weights provide for relatively higher amounts of bandwidth.

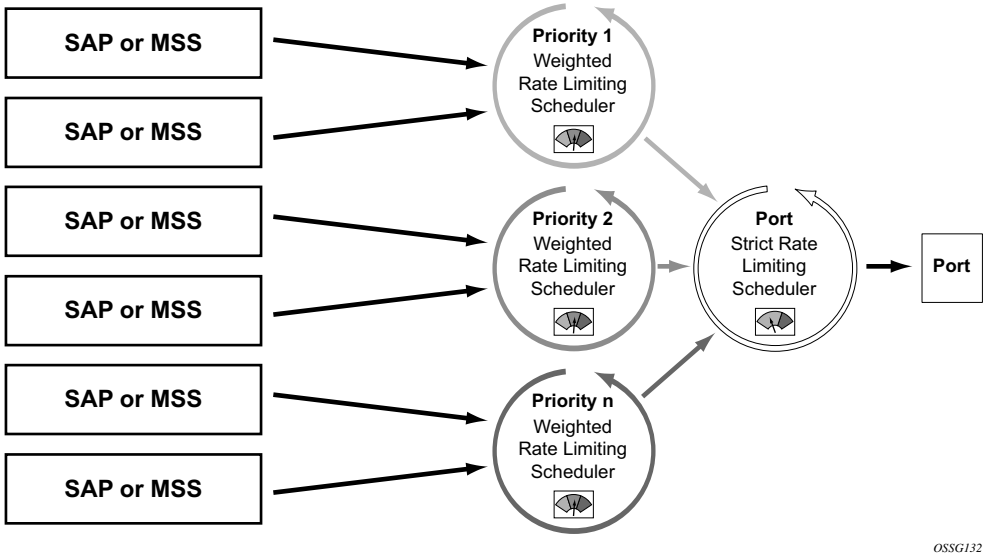

Figure 19: Schedulers on SAP or Multi-Service Site Receive Bandwidth From Port Priority Levels

Direct Service or Subscriber Queue Association to Port Scheduler Parents

The second model of bandwidth allocation on an egress access port is to directly associate a service or subscriber queue to a port-level scheduler. This model allows the port scheduler hierarchy to allocate bandwidth on a per class or priority basis to each service or subscriber queue. This allows the provider to manage the available egress port bandwidth on a service tier basis ensuring that during egress port congestion, a deterministic behavior is possible from an aggregate perspective. While this provides an aggregate bandwidth allocation model, it does not inhibit per service or per subscriber queuing. Figure 20 demonstrates the single, port scheduler policy model.

Figure 20 also demonstrates the optional aggregate rate limiter at the SAP, multi-service site or subscriber level. The aggregate rate limiter is used to define a maximum aggregate bandwidth at which the child queues can operate. While the port-level scheduler is allocating bandwidth to each child queue, the current sum of the bandwidth for the service or subscriber is monitored. Once the aggregate rate limit is reached, no more bandwidth is allocated to the children associated with the SAP, multi-service site, or subscriber. Aggregate rate limiting is restricted to the single scheduler policy model and is mutually exclusive to defining SAP, multi-service site, or subscriber scheduling policies.

The benefit of the single scheduler policy model is that the bandwidth is allocated per priority for all queues associated with the egress port. This allows a provider to preferentially allocate bandwidth to higher priority classes of service independent of service or subscriber instance. In many cases, a subscriber can purchase multiple services from a single site (VoIP, HSI, Video) and each service can have a higher premium value relative to other service types. If a subscriber has purchased a premium service class, that service class should get bandwidth before another subscriber's best effort service class. When combined with the aggregate rate limit feature, the single port-level scheduler policy model provides a per-service instance or per-subscriber instance aggregate SLA and a class based port bandwidth allocation function.

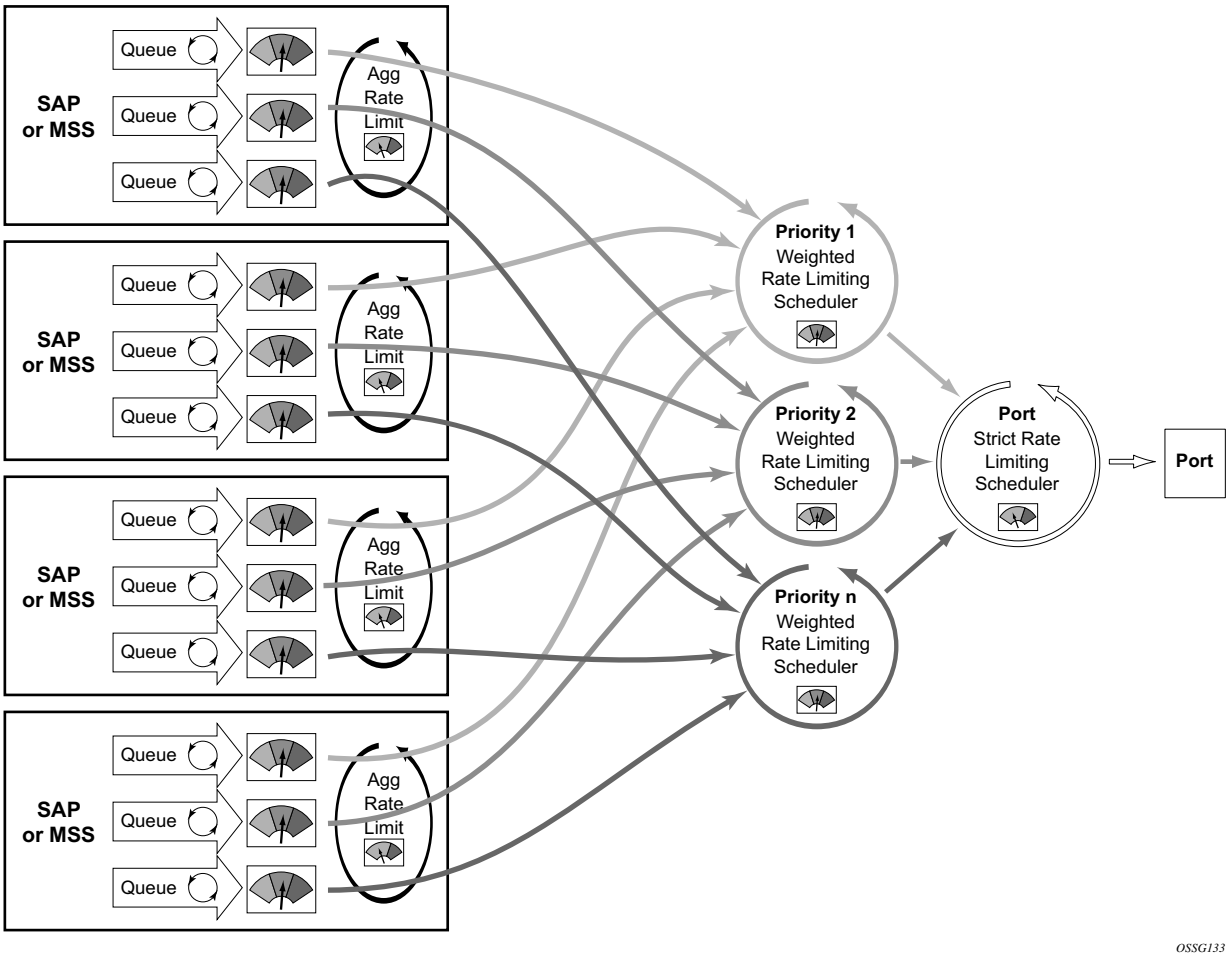

Figure 20: Direct Service or Subscriber Association to Port Scheduler Model

Frame and Packet-Based Bandwidth Allocation

A port-based bandwidth allocation mechanism must consider the effect that line encapsulation overhead plays relative to the bandwidth allocated per service or subscriber. The service or subscriber level bandwidth definition (at the queue level) operates on a packet accounting basis. For Ethernet, this includes the DLC header, the payload and the trailing CRC. This does not include the IFG or the preamble. This means that an Ethernet packet will consume 20 bytes more bandwidth on the wire than what the queue accounted for. When considering HDLC encoded PoS or SDH ports, the overhead is variable based on '7e' insertions (and other TDM framing issues). The HDLC and SONET/SDH frame overhead is not included for queues forwarding on PoS and SDH links.

The port-based scheduler hierarchy must translate the frame based accounting (on-the-wire bandwidth allocation) it performs to the packet based accounting in the queues. When the port scheduler considers the maximum amount of bandwidth a queue should get, it must first determine how much bandwidth the queue can use. This is based on the offered load the queue is currently experiencing (how many octets are being offered the queue). The offered load is compared to the queues configured CIR and PIR. The CIR value determines how much of the offered load should be considered in the "within-cir" bandwidth allocation pass. The PIR value determines how much of the remaining offered load (after "within-cir") should be considered for the "above-cir" bandwidth allocation pass.

For Ethernet queues (queues associated with an egress Ethernet port), the packet to frame conversion is relatively easy. The system multiplies the number of offered packets by 20 bytes and adds the result to the offered octets (offeredPackets x 20 + offeredOctets = frameOfferedLoad). This frame-offered-load value represents the amount of line rate bandwidth the queue is requesting. The system computes the ratio of increase between the offered-load and frame-offered-load and calculates the current frame based CIR and PIR. The frame-CIR and frame-PIR values are used as the limiting values in the "within-cir" and "above-cir" port bandwidth distribution passes.

For PoS or SDH queues, the packet to frame conversion is more difficult to dynamically calculate due to the variable nature of HDLC encoding. Wherever a '7e' bit or byte pattern appears in the data stream, the framer performing the HDLC encoding must place another '7e' within the payload. Since this added HDLC encoding is unknown to the forwarding plane, the system allows for an encapsulation overhead parameter that can be provisioned on a per queue basis. This is provided on a per queue basis to allow for differences in the encapsulation behavior between service flows in different queues. The system multiplies the offered load of the queue by the encapsulation-overhead parameter and adds the result to the offered load of the queue (offeredOctets * configuredEncapsulationOverhead + offeredOctets = frameOfferedLoad). The frame-offered-load value is used by the egress PoS/SDH port scheduler in the same manner as the egress Ethernet port scheduler above.

From a provisioning perspective, queues and service level (and subscriber level) scheduler policies are always provisioned with packet-based parameters. The system will convert these values to

frame-based on-the-wire values for the purpose of port bandwidth allocation. However, port-based scheduler policy scheduler maximum rates and CIR values are always interpreted as on-the-wire values and must be provisioned accordingly. Figure 21 and Figure 22 provide a logical view of bandwidth distribution from the port to the queue level and shows the packet or frame-based provisioning at each step.

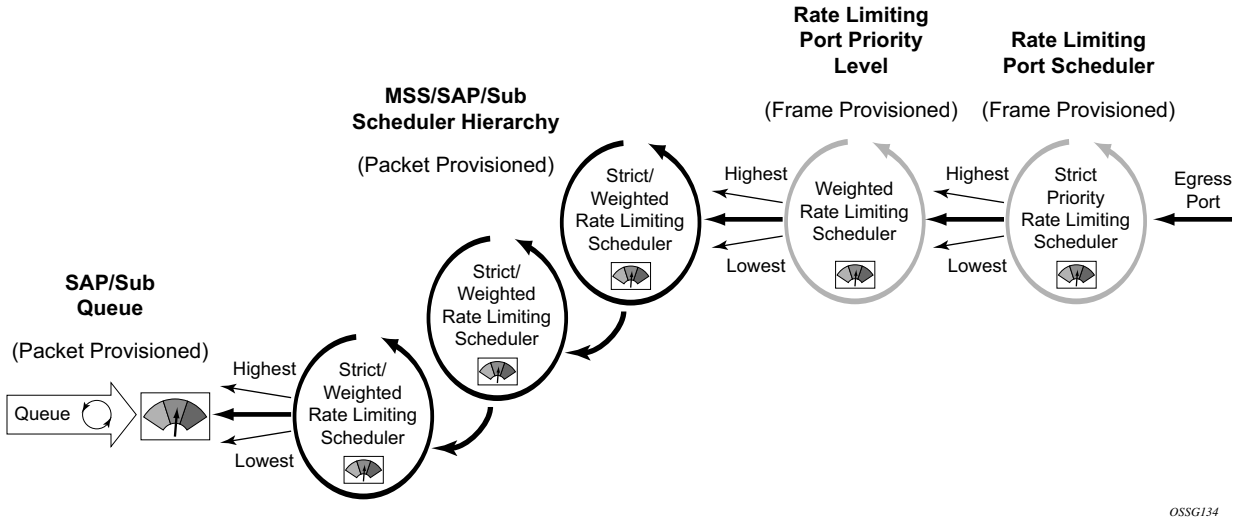

Figure 21: Port Bandwidth Distribution for Service and Port Scheduler Hierarchies

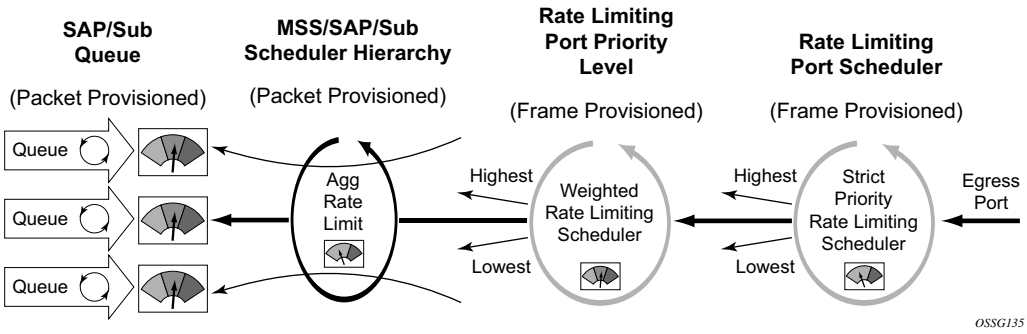

Figure 22: Port Bandwidth Distribution for Direct Queue to Port Scheduler Hierarchy

Queue Parental Association Scope

A **port-parent** command in the sap-egress and network-queue QoS policy queue context defines the direct child/parent association between an egress queue and a port scheduler priority level. The **port-parent** command is mutually exclusive to the already-existing **parent** command, which associates a queue with a scheduler at the SAP, multi-service site or subscriber profile level. It is possible to mix local parented (parent to service or subscriber level scheduler) and port parented queues with schedulers on the same egress port.

The **port-parent** command only accepts a child/parent association to the eight priority levels on a port scheduler hierarchy. Similar to the local **parent** command, two associations are supported, one for "within-cir" bandwidth (cir-level) and a second one for "above-cir" bandwidth (level). The "within-cir" association is optional and can be disabled by using the default "within-cir" weight value of 0. In the event that a queue with a defined parent port is on a port without a port scheduler policy applied, that queue will be considered an orphaned queue. If a queue with a parent command is defined on a port and the named scheduler is not found due a missing scheduler policy or a missing scheduler of that name, the queue will be considered orphaned as well.

A queue can be moved from a local (on the SAP, multi-service site, or subscriber profile) parent to a port parent priority level simply by executing the **port-parent** command. Once the **port-parent** command is executed, any local parent information for the queue is lost. The queue can also be moved back to a local parent at anytime by executing the local parent command. Lastly, the local parent or port parent association can be removed at any time by using the no version of the appropriate parent command.

Service or Subscriber-Level Scheduler Parental Association Scope

The **port-parent** command in the scheduler-policy scheduler context (at all tier levels) allows a scheduler to be associated with a port scheduler priority level. The **port-parent** command is mutually exclusive to the **parent** command for schedulers at tiers 2 and 3 within the scheduler policy. The **port-parent** command is the only parent command allowed for schedulers in tier 1.

The **port-parent** command only accepts a child/parent association to the eight priority levels on a port scheduler hierarchy. Similar to the normal local parent command, two associations are supported, one for "within-cir" bandwidth (cir-level) and a second one for "above-cir" bandwidth (level). The "within-cir" association is optional and can be disabled by using the default "within-cir" weight value of 0. In the event that a scheduler with a port parent defined is on a port without a port scheduler policy applied, that scheduler will be considered an orphaned scheduler.

A scheduler in tiers 2 and 3 can be moved from a local (within the policy) parent to a port parent priority level simply by executing the **port-parent** command. Once the **port-parent** command is executed, any local parent information for the scheduler is lost. The schedulers at tiers 2 and 3 can also be moved back to a local parent at anytime by executing the local parent command. Lastly, the

local parent or port parent association can be removed at anytime by using the no version of the appropriate parent command. A scheduler in tier 1 can only be associated with a port parent and that port parent definition can be added or removed at anytime.

Network Queue Parent Scheduler

Network queues support port scheduler parent priority-level associations. Using a port scheduler policy definition and mapping network queues to a port parent priority level, HQoS functionality is supported providing eight levels of strict priority and weights within the same priority. A network queue's bandwidth is allocated using the "within-cir" and "above-cir" scheme normal for port schedulers.

Queue CIR and PIR percentages when port-based schedulers are in effect will be based on frame-offered-load calculations. Figure 23 demonstrates port-based virtual scheduling bandwidth distribution.

A network queue with a port parent association exists on a port without a scheduler policy defined will be considered to be orphaned.

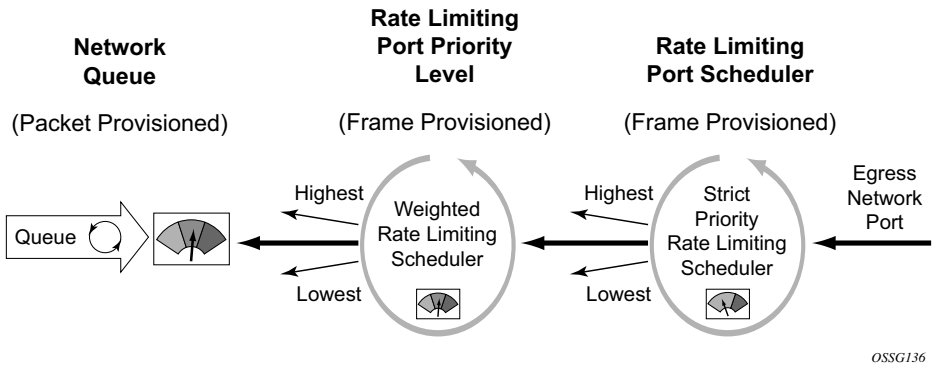

Figure 23: Bandwidth Distribution on Network Port with Port-Based Scheduling

Foster Parent Behavior for Orphaned Queues and Schedulers

All queues and schedulers on a port that has a port-based scheduler policy configured will be subject to bandwidth allocation through the port-based schedulers. All queues and schedulers that are not configured with a scheduler parent are considered to be orphaned when port-based scheduling is in effect. This includes access and network queue schedulers at the SAP, multi-service site, subscriber and port level.

By default, orphaned queues and schedulers are allocated bandwidth after all queues and schedulers in the parented hierarchy have had bandwidth allocated “within-cir” and “above-cir”. In essence, an orphaned scheduler or queue can be considered as being foster parented by the port scheduler. Orphaned queues and schedulers have an inherent port scheduler association as shown below:

- Within-CIR priority = 1
- Within-CIR weight = 0
- Above-CIR priority = 1
- Above-CIR weight = 0

The above-CIR weight = 0 value is only used for orphaned queues and schedulers on port scheduler enabled egress ports. The system interprets weight=0 as priority level 0 and will only distribute bandwidth to level 0 once all other properly parented queues and schedulers have received bandwidth. Orphaned queues and schedulers all have equal priority to the remaining port bandwidth.

The default orphan behavior can be overridden for each port scheduler policy by using the orphan override command. The orphan override command accepts the same parameters as the port parent command. When the orphan override command is executed, all orphan queues and schedulers are treated in a similar fashion as other properly parented queues and schedulers based on the override parenting parameters.

It is expected that an orphan condition is not the desired state for a queue or scheduler and is the result of a temporary configuration change or configuration error.

Frame-Based Accounting

The standard accounting mechanism uses ‘packet based’ rules that account for the DLC header, any existing tags, Ethernet payload and the 4 byte CRC. The Ethernet framing overhead which includes the Inter-Frame Gap (IFG) and preamble (20 bytes total) are not included in packet based accounting. When frame based accounting is enabled, the 20 byte framing overhead is included in the queue CIR, PIR and scheduling operations allowing the operations to take into consideration on-wire bandwidth consumed by each Ethernet packet.

Since the native queue accounting functions (stats, CIR and PIR) are based on packet sizes and do not include Ethernet frame encapsulation overhead, the system must manage the conversion between packet based and frame based accounting. To accomplish this, the system requires that a queue operates in frame based accounting mode, and must be managed by a virtual scheduler policy or by a port virtual scheduler policy. Egress queues can use either port or service schedulers to accomplish frame based accounting, but ingress queues are limited to service based scheduling policies.

Turning on frame based accounting for a queue is accomplished through a frame based accounting command defined on the scheduling policy level associated with the queue or through a queue frame based accounting parameter on the aggregate rate limit command associated with the queues SAP, multi-service site or subscriber context.

Operational Modifications

To add frame overhead to the existing QoS Ethernet packet handling functions, the system uses the already existing virtual scheduling capability of the system. The system currently monitors each queue included in a virtual scheduler to determine its offered load. This offered load value is interpreted based on the queues defined CIR and PIR threshold rates to determine bandwidth offerings from the queues virtual scheduler. When egress port based virtual scheduling was added, frame based usage on the wire was added to allow for the port bandwidth to be accurately allocated to each child queue on the port.

Existing Egress Port Based Virtual Scheduling

The port based virtual scheduling mechanism takes the native packet based accounting results from the queue and adds 20 bytes to each packet to derive the queue's frame based offered load. The ratio between the frame based offered load and the packet based offered load is then used to determine the effective frame based CIR and frame based PIR thresholds for the queue. Once the port virtual scheduler computes the amount of bandwidth allowed to the queue (in a frame based fashion), the bandwidth is converted back to a packet based value and used as the queue's operational PIR. The queue's native packet based mechanisms continue to function, but the maximum operational rate is governed by frame based decisions.

Queue Behavior Modifications for Frame Based Accounting

The frame based accounting feature extends this capability to allow the queue CIR and PIR thresholds to be defined as frame based values as opposed to packet based values. The queue continues to internally use its packet based mechanisms, but the provisioned frame based CIR and PIR values are continuously revalued based on the ratio between the calculated frame based offered load and actual packet based offered load. As a result, the queue's operational packet based CIR and PIR are accurately modified during each iteration of the virtual scheduler to represent the provisioned frame based CIR and PIR.

Virtual Scheduler Rate and Queue Rate Parameter Interpretation

Normally, a scheduler policy contains rates that indicate packet based accounting values. When the children queues associated with the policy are operating in frame based accounting mode, the parent schedulers must also be governed by frame based rates. Since either port based or service based virtual scheduling is required for queue frame based operation, enabling frame based operation is configured at either the scheduling policy or aggregate rate limit command level. All queues associated with the policy or the aggregate rate limit command will inherit the frame based accounting setting from the scheduling context.

When frame based accounting is enabled, the queues CIR and PIR settings are automatically interpreted as frame based values. If a SAP ingress QoS policy is applied with a queue PIR set to 100Mbps on two different SAPs, one associated with a policy with frame based accounting enabled and the other without frame based accounting enabled, the 100Mbps rate will be interpreted differently for each queue. The frame based accounting queue will add 20 bytes to each packet received by the queue and limit the rate based on the extra overhead. The packet based accounting queue will not add the 20 bytes per packet and thus allow more packets through per second.

Similarly, the rates defined in the scheduling policy with frame based accounting enabled will automatically be interpreted as frame based rates.

The port based scheduler aggregate rate limit command always interprets its configured rate limit value as a frame based rate. Setting the frame based accounting parameter on the aggregate rate limit command only affects the queues managed by the aggregate rate limit and converts them from packet based to frame based accounting mode.

Configuring Port Scheduler Policies

Port Scheduler Structure

Every port scheduler supports eight strict priority levels with a two pass bandwidth allocation mechanism for each priority level. Priority levels 8 through 1 (level 8 is the highest priority) are available for port-parent association for child queues and schedulers. Each priority level supports a maximum rate limit parameter that limits the amount of bandwidth that may be allocated to that level. A CIR parameter is also supported that limits the amount of bandwidth allocated to the priority level for the child queue's offered load, within their defined CIR. An overall maximum rate parameter defines the total bandwidth that will be allocated to all priority levels.

Special Orphan Queue and Scheduler Behavior

When a port scheduler is present on an egress port or channel, the system ensures that all queues and schedulers receive bandwidth from that scheduler to prevent free-running queues which can cause the aggregate operational PIR of the port or channel to oversubscribe the bandwidth available. When the aggregate maximum rate for the queues on a port or channel operate above the available line rate, the forwarding ratio between the queues will be affected by the hardware schedulers on the port and may not reflect the scheduling defined on the port or intermediate schedulers. Queues and schedulers that are either explicitly attached to the port scheduler using the port-parent command or are attached to an intermediate scheduler hierarchy that is ultimately attached to the port scheduler are managed through the normal eight priority levels. Queues and schedulers that are not attached directly to the port scheduler and are not attached to an intermediate scheduler that itself is attached to the port scheduler are considered orphaned queues and, by default, are tied to priority 1 with a weight of 0. All weight 0 queues and schedulers at priority level 1 are allocated bandwidth after all other children and each weight 0 child is given an equal share of the remaining bandwidth. This default orphan behavior may be overridden at the port scheduler policy by using the orphan-override command. The orphan-override command accepts the same parameters as the port-parent command. When the orphan-override command is executed, the parameters will be used as the port parent parameters for all orphans associated with a port using the port scheduler policy.

Packet to Frame Bandwidth Conversion

Another difference between the service level scheduler-policy and the port level port-scheduler-policy is in bandwidth allocation behavior. The port scheduler is designed to offer on-the-wire bandwidth. For Ethernet ports, this includes the IFG and the preamble for each frame and represents 20 bytes total per frame. The queues and intermediate service level schedulers (a

service level scheduler is a scheduler instance at the SAP, multi-service site or subscriber profile level) operate based on packet overhead which does not include the IFG or preamble on Ethernet packets. In order for the port based virtual scheduling algorithm to function, it must convert the queue and service scheduler packet based required bandwidth and bandwidth limiters (CIR and rate PIR) to frame based values. This is accomplished by adding 20 bytes to each Ethernet frame offered at the queue level to calculate a frame based offered load. Then the algorithm calculates the ratio increase between the packet based offered load and the frame based offered load and uses this ratio to adapt the CIR and rate PIR values for the queue to frame-CIR and frame-PIR values. When a service level scheduler hierarchy is between the queues and the port based schedulers, the ratio between the average frame-offered-load and the average packet-offered-load is used to adapt the scheduler's packet based CIR and rate PIR to frame based values. The frame based values are then used to distribute the port based bandwidth down to the queue level.

Packet over SONET (PoS) and SDH queues also operate based on packet sizes and do not include on-the-wire frame overhead. Unfortunately, the port based virtual scheduler algorithm does not have access to all the frame encapsulation overhead occurring at the framer level. Instead of automatically calculating the difference between packet-offered-load and frame-offered-load, the system relies on a provisioned value at the queue level. This avg-frame-overhead parameter is used to calculate the difference between the packet-offered-load and the frame-offered-load. This difference is added to the packet-offered-load to derive the frame-offered-load. Proper setting of this percentage value is required for proper bandwidth allocation between queues and service schedulers. If this value is not attainable, another approach is to artificially lower the maximum rate of the port scheduler to represent the average port framing overhead. This, in conjunction with a zero or low value for avg-frame-overhead, will ensure that the allocated queue bandwidth will control forwarding behavior instead of the low level hardware schedulers.

Aggregate Rate Limits for Directly Attached Queues

When all queues for a SAP, multi-service site or subscriber instance are attached directly to the port scheduler (using the port-parent command), it is possible to configure an agg-rate-limit for the queues. This is beneficial since the port scheduler does not provide a mechanism to enforce an aggregate SLA for a service or subscriber and the agg-rate-limit provides this ability. Queues may be provisioned directly on the port scheduler when it is desirable to manage the congestion at the egress port based on class priority instead of on a per service object basis.

The agg-rate-limit is not supported when one or more queues on the object are attached to an intermediate service scheduler. In this event, it is expected that the intermediate scheduler hierarchy will be used to enforce the aggregate SLA. Attaching an agg-rate-limit is mutually exclusive to attaching an egress scheduler policy at the SAP or subscriber profile level. Once an aggregate rate limit is in effect, a scheduler policy cannot be assigned. Once a scheduler policy is assigned on the egress side of a SAP or subscriber profile, an agg-rate-limit cannot be assigned.

Since the sap-egress policy defines a queue's parent association before the policy is associated with a service SAP or subscriber profile, it is possible for the policy to either not define a port-parent association or define an intermediate scheduler parenting that does not exist. As stated above, queues in this state are considered to be orphaned and automatically attached to port scheduler priority 1. Orphaned queues are included in the aggregate rate limiting behavior on the SAP or subscriber instance they are created within.

SAP Egress QoS Policy Queue Parenting

A sap-egress QoS policy queue may be associated with either a port parent or an intermediate scheduler parent. The validity parent definition cannot be checked at the time it is provisioned since the application of the QoS policy is not known until it is applied to an egress SAP or subscriber profile. It is allowed to have port or intermediate parenting decided on a queue by queue basis, some queues tied directly to the port scheduler priorities while other queues are attached to intermediate schedulers.

Network Queue QoS Policy Queue Parenting

A network-queue policy only supports direct port parent priority association. Intermediate schedulers are not supported on network ports or channels.

Egress Port Scheduler Overrides

Once a port scheduler has been associated with an egress port, it is possible to override the following parameters:

- The max-rate allowed for the scheduler.
- The maximum rate for each priority level 8 through 1.
- The CIR associated with each priority level 8 through 1.

The orphan priority level (level 1) has no configuration parameters and cannot be overridden.

Applying a Port Scheduler Policy to a Virtual Port

In order to represent a downstream network aggregation node in the local node scheduling hierarchy, a new scheduling node, referred to as virtual port, and vport in CLI have been introduced. The vport operates exactly like a port scheduler except multiple vport objects can be configured on the egress context of an Ethernet port.

Figure 24 illustrates the use of the vport on an Ethernet port of a Broadband Network Gateway (BNG). In this case, the vport represents a specific downstream DSLAM.

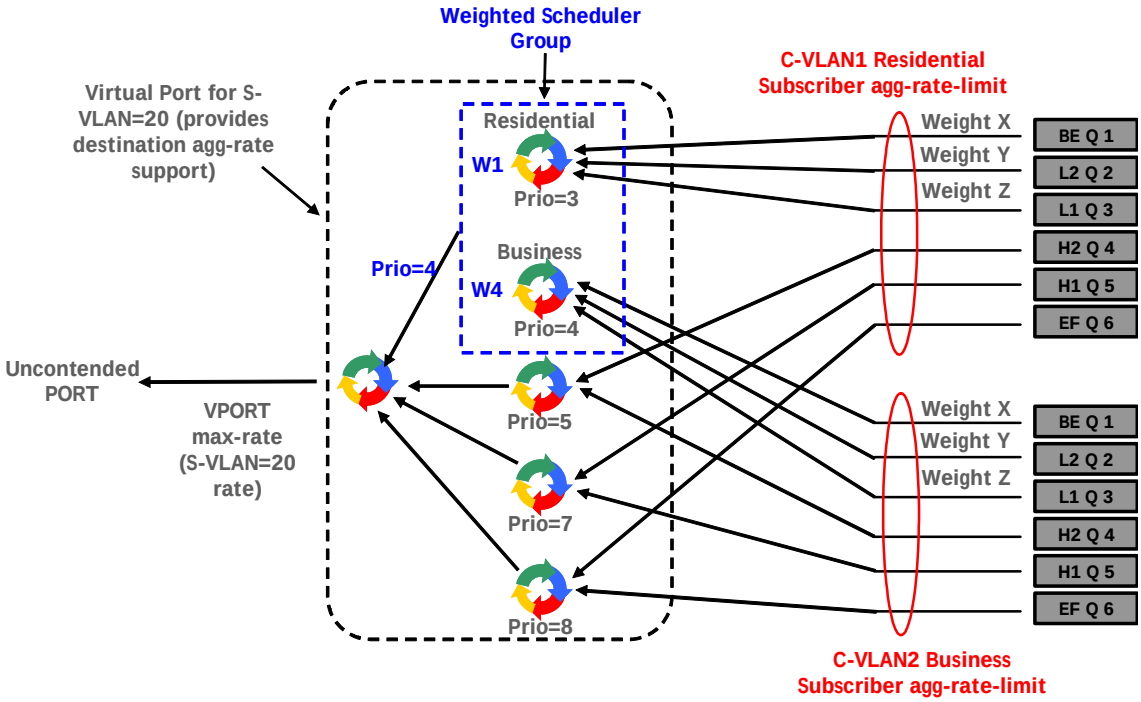

Figure 24: Applying a Port Scheduler Policy to a Vport

The user adds a vport to an Ethernet port using the following command:

CLI Syntax: `configure>port>ethernet>access>egress>vport` *`vport-name`* `create`

The vport is always configured at the port level even when a port is a member of a LAG. The vport name is local to the port it is applied to but must be the same for all member ports of a LAG. It however does not need to be unique globally on a chassis.

The user applies a port scheduler policy to a vport using the following command:

CLI Syntax: `configure>port>ethernet>acess>egress>vport>port-scheduler-policy` *`port-scheduler-policy-name`*

A Vport cannot be parented to the port scheduler when it is using a port scheduler policy itself. It is thus important the user ensures that the sum of the **max-rate** parameter value in the port scheduler policies of all Vport instances on a given egress Ethernet port does not oversubscribe the port's rate. If it does, the scheduling behavior degenerates to that of the H/W scheduler on that port. A Vport which uses an **agg-rate,** or a scheduler-policy can be parented to a port scheduler. This is explained in Section Applying Aggregate Rate Limit to a VPORT. Note that the application of the **agg-rate** *rate*, **port-scheduler-policy** and **scheduler-policy** commands under a VPORT are mutually exclusive.

Each subscriber host queue is port parented to the Vport which corresponds to the destination DSLAM using the existing port-parent command:

CLI Syntax: `configure>qos>sap-egress>queue>port-parent [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]`

This command can parent the queue to either a port or to a Vport. These operations are mutually exclusive in CLI as explained above. When parenting to a Vport, the parent Vport for a subscriber host queue is not explicitly indicated in the above command. It is determined indirectly. The determination of the parent Vport for a given subscriber host queue is described in the 7750 SR OS Triple Play Guide.

Subscriber host queues, SLA profile schedulers, subscriber profile scheduler and PW SAPs (in IES or VPRN services) can be parented to a VPORT.

Applying Aggregate Rate Limit to a VPORT

The user can apply an aggregate rate limit to the Vport and apply a port scheduler policy to the port.

This model allows the user to over subscribe the Ethernet port. The application of the **agg-rate** option is mutually exclusive with the application of a port scheduler policy, or a scheduler policy to a VPORT.

When using this model, a subscriber host queue with the **port-parent** option enabled is scheduled within the context of the port's port scheduler policy. More details are provided in the 7750 SR OS Triple Play Guide.

Applying a Scheduler Policy to a VPORT

The user can apply a scheduler policy to the VPORT. This is allows scheduling control of subscriber tier 1 schedulers in a scheduler policy applied to the egress of a subscriber or SLA profile, or to a PW SAP in an IES or VPRN service.

The advantage of using a scheduler policy under a VPORT, compared to the use of a port scheduler (with or without an agg-rate rate), is that it allows a port parent to be configured at the VPORT level as well as allowing the user to oversubscribe the Ethernet port.

Bandwidth distribution from an egress port scheduler to a VPORT configured with a scheduler policy can be performed based on the level/cir-level and weight/cir-weight configured under the scheduler's port parent. The result is in allowing multiple VPORTs, for example representing different DSLAMs, to share the port bandwidth capacity in a flexible way that is under the control of the user.

The configuration of a scheduler policy under a VPORT is mutually exclusive with the configuration of a port scheduler policy or an aggregate rate limit.

A scheduler policy is configured under a VPORT as follows:

CLI Syntax: config>port>ethernet>access>egress>vport# **scheduler-policy** *scheduler-policy-name*

When using this model, a tier 1 scheduler in a scheduling policy applied to a subscriber profile or SLA profiles must be configured as follows:

CLI Syntax: config>qos>scheduler-policy>tier# **parent-location vport**

If the VPORT exists, but the port does not have a port scheduler policy applied, then its schedulers will be orphaned and no port level QOS control can be enforced.

The following **show**/**clear** commands are available related to the VPORT scheduler:

```
show qos scheduler-hierarchy port port-id vport name [scheduler scheduler-name] [detail]

show qos scheduler-stats port port-id vport name [scheduler scheduler-name] [detail]

clear qos scheduler-stats port port-id vport name [scheduler scheduler-name] [detail]
```

HQoS adjustment and host tracking are not supported on schedulers that are configured in a scheduler policy on a VPORT, so the configuration of a scheduler policy under a VPORT is mutually exclusive with the configuration of the egress-rate-modify parameter.

ESM over MPLS pseudowires are not supported when a scheduler policy is configured on a VPORT.

Weighted Scheduler Group in a Port Scheduler Policy

The existing port scheduler policy defines a set of eight priority levels with no ability of grouping levels within a single priority. In order to allow for the application of a scheduling weight to groups of queues competing at the same priority level of the port scheduler policy applied to the vport, or to the Ethernet port, a new group object is defined under the port scheduler policy:

CLI Syntax: `configure>qos>port-scheduler-policy>group group-name rate pir-rate[cir cir-rate]`

Up to eight groups can be defined within each port scheduler policy. One or more levels can map to the same group. A group has a rate and optionally a cir-rate and inherits the highest scheduling priority of its member levels. For example, the scheduler group shown in the vport in Figure 24 consists of level priority 3 and level priority 4. It thus inherits priority 4 when competing for bandwidth with the standalone priority levels 8, 7, and 5.

In essence, a group receives bandwidth from the port or from the vport and distributes it within the member levels of the group according to the weight of each level within the group. Each priority level will compete for bandwidth within the group based on its weight under congestion situation. If there is no congestion, a priority level can achieve up to its rate (cir-rate) worth of bandwidth.

The mapping of a level to a group is performed as follows:

CLI Syntax: `configure>qos>port-scheduler-policy>level priority-level rate pir-rate [cir cir-rate]  group group-name [weight weight-in-group]`

Note that CLI will enforce that mapping of levels to a group are contiguous. In other words, a user would not be able to add priority level to group unless the resulting set of priority levels is contiguous.

When a level is not explicitly mapped to any group, it maps directly to the root of the port scheduler at its own priority like in existing behavior.

Basic Configurations

A basic QoS scheduler policy must conform to the following:

- Each QoS scheduler policy must have a unique policy ID.
- A tier level 1 parent scheduler name cannot be configured.

A basic QoS port scheduler policy must conform to the following:

- Each QoS port scheduler policy must have a unique policy name.

Create a QoS Scheduler Policy

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP or IP interface, a default QoS policy is applied.

To create a scheduler policy, define the following:

- A scheduler policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Specify the tier level. A tier identifies the level of hierarchy that a group of schedulers are associated with.
- Specify a scheduler name. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler.
- Specify a parent scheduler name to be associated with a level 2 or 3 tier.
- You can modify the bandwidth that the scheduler can offer its child queues or schedulers. Otherwise, the scheduler will be allowed to consume bandwidth without a scheduler-defined limit.

The following displays a scheduler policy configuration:

```
A:ALA-12>config>qos# info
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
     scheduler-policy "SLA1" create
           description "NetworkControl(3), Voice(2) and NonVoice(1) have strict priorities"
            tier 1
                scheduler "All_traffic" create
                    description "All traffic goes to this scheduler eventually"
                    rate 11000
                exit
            exit
            tier 2
```

```
                scheduler "NetworkControl" create
                    description "network control traffic within the VPN"
                    parent All_traffic level 3 cir-level 3
                    rate 100
                exit
                scheduler "NonVoice" create
                    description "NonVoice of VPN and Internet traffic will be serviced by
this scheduler"
                    parent All_traffic cir-level 1
                    rate 11000
                exit
                scheduler "Voice" create
                   description "Any voice traffic from VPN and Internet use this scheduler"
                    parent All_traffic level 2 cir-level 2
                    rate 5500
                exit
            exit
            tier 3
                scheduler "Internet_be" create
                    parent NonVoice cir-level 1
                exit
                scheduler "Internet_priority" create
                    parent NonVoice level 2 cir-level 2
                exit
                scheduler "Internet_voice" create
                    parent Voice
                exit
                scheduler "VPN_be" create
                    parent NonVoice cir-level 1
                exit
                scheduler "VPN_nc" create
                    parent NetworkControl
                    rate 100 cir 36
                exit
                scheduler "VPN_priority" create
                    parent NonVoice level 2 cir-level 2
                exit
                scheduler "VPN_reserved" create
                    parent NonVoice level 3 cir-level 3
                exit
                scheduler "VPN_video" create
                    parent NonVoice level 5 cir-level 5
                    rate 1500 cir 1500
                exit
                scheduler "VPN_voice" create
                    parent Voice
                    rate 2500 cir 2500
                exit
            exit
        exit
        sap-ingress 100 create
            description "Used on VPN sap"
...
-----------------------------------------------
A:ALA-12>config>qos#
```

Applying Scheduler Policies

Apply scheduler policies to the following entities:

- Customer
- Epipe
- IES
- VPLS
- VPRN

Customer

Use the following CLI syntax to associate a scheduler policy to a customer's multiservice site:

CLI Syntax:
```
config>customer customer-id
 multiservice-site customer-site-name
    egress
       scheduler-policy scheduler-policy-name
    ingress
       scheduler-policy scheduler-policy-name
```

Epipe

Use the following CLI syntax to apply QoS policies to ingress and/or egress Epipe SAPs:

CLI Syntax:
```
config>service# epipe service-id [customer customer-id]
    sap sap-id
     egress
        scheduler-policy scheduler-policy-name
     ingress
        scheduler-policy scheduler-policy-name
```

CLI Syntax:
```
config>service# epipe service-id [customer customer-id]
    sap sap-id
     egress
        qos sap-egress-policy-id
     ingress
        qos sap-ingress-policy-id
```

The following output displays an Epipe service configuration with SAP scheduler policy SLA2 applied to the SAP ingress and egress.

```
A:SR>config>service# info
----------------------------------------------
        epipe 6 customer 6 vpn 6 create
            description "Distributed Epipe service to west coast"
            sap 1/1/10:0 create
                ingress
                    scheduler-policy "SLA2"
                    qos 100
                exit
                egress
                    scheduler-policy "SLA2"
                    qos 1010
                exit
            exit
...
----------------------------------------------
A:SR>config>service#
```

IES

Use the following CLI syntax to apply scheduler policies to ingress and/or egress IES SAPs:

CLI Syntax: config>service# ies *service-id* [customer *customer-id*]
interface *ip-int-name*
sap *sap-id*
egress
scheduler-policy *scheduler-policy-name*
ingress
scheduler-policy *scheduler-policy-name*

The following output displays an IES service configuration with scheduler policy SLA2 applied to the SAP ingress and egress.

```
A:SR>config>service# info
----------------------------------------------
        ies 88 customer 8 vpn 88 create
            interface "Sector A" create
                sap 1/1/1.2.2 create
                    ingress
                        scheduler-policy "SLA2"
                        qos 101
                    exit
                    egress
                        scheduler-policy "SLA2"
                        qos 1020
                    exit
                exit
            exit
            no shutdown
        exit
----------------------------------------------
A:SR>config>service#
```

VPLS

Use the following CLI syntax to apply scheduler policies to ingress and/or egress VPLS SAPs:

CLI Syntax:
```
config>service# vpls service-id [customer customer-id]
                sap sap-id
                 egress
                    scheduler-policy scheduler-policy-name
                 ingress
                    scheduler-policy scheduler-policy-name
```

The following output displays an VPLS service configuration with scheduler policy SLA2 applied to the SAP ingress and egress.

```
A:SR>config>service# info
----------------------------------------------
...
        vpls 700 customer 7 vpn 700 create
            description "test"
            stp
                shutdown
            exit
            sap 1/1/9:0 create
                ingress
                    scheduler-policy "SLA2"
                    qos 100
                exit
                egress
                    scheduler-policy "SLA2"
                exit
            exit
            spoke-sdp 2:222 create
            exit
            mesh-sdp 2:700 create
            exit
            no shutdown
        exit
...
----------------------------------------------
A:SR>config>service#
```

VPRN

Use the following CLI syntax to apply scheduler policies to ingress and/or egress VPRN SAPs:

CLI Syntax:
```
config>service# vprn service-id [customer customer-id]
                interface ip-int-name
                   sap sap-id
                      egress
                         scheduler-policy scheduler-policy-name
                      ingress
                         scheduler-policy scheduler-policy-name
```

The following output displays a VPRN service configuration with the scheduler policy SLA2 applied to the SAP ingress and egress.

```
A:SR7>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            ecmp 8
            autonomous-system 10000
            route-distinguisher 10001:1
            auto-bind ldp
            vrf-target target:10001:1
            interface "to-ce1" create
                address 11.1.0.1/24
                sap 1/1/10:1 create
                    ingress
                        scheduler-policy "SLA2"
                    exit
                    egress
                        scheduler-policy "SLA2"
                    exit
                exit
            exit
            no shutdown
        exit
        epipe 6 customer 6 vpn 6 create
----------------------------------------------
A:SR7>config>service#
```

Creating a QoS Port Scheduler Policy

Configuring and applying QoS port scheduler policies is optional. If no QoS port scheduler policy is explicitly applied to a SAP or IP interface, a default QoS policy is applied.

To create a port scheduler policy, define the following:

- A port scheduler policy name.
- Include a description. The description provides a brief overview of policy features.

Use the following CLI syntax to create a QoS port scheduler policy.

Note that the **create** keyword is included in the command syntax upon creation of a policy.

CLI Syntax:
```
config>qos
 port-scheduler-policy scheduler-policy-name [create]
    description description-string
    level priority-level rate pir-rate [cir cir-rate]
    max-rate rate
    orphan-override [level priority-level] [weight weight]
       [cir-level priority-level] [cir-weight cir-weight]
```

The following displays a scheduler policy configuration example:

```
*A:ALA-48>config>qos>port-sched-plcy# info
----------------------------------------------
            description "Test Port Scheduler Policy"
            orphan-override weight 50 cir-level 4 cir-weight 50
----------------------------------------------
*A:ALA-48>config>qos>port-sched-plcy#
```

Configuring Port Parent Parameters

The **port-parent** command defines a child/parent association between an egress queue and a port based scheduler or between an intermediate service scheduler and a port based scheduler. The command may be issued in three distinct contexts; **sap-egress>queue** *queue-id*, and **network-queue> queue** *queue-id* and **scheduler-policy>scheduler** *scheduler-name* the **network-queue> queue** *queue-id* context. The **port-parent** command allows for a set of within-cir and above-cir parameters that define the port priority levels and weights for the queue or scheduler. If the port-parent command is executed without any parameters, the default parameters are assumed.

Within-CIR Priority Level Parameters

The within-cir parameters define which port priority level the queue or scheduler should be associated with when receiving bandwidth for the queue or schedulers within-cir offered load. The within-cir offered load is the amount of bandwidth the queue or schedulers could use that is equal to or less than its defined or summed CIR value. The summed value is only valid on schedulers and is the sum of the within-cir offered loads of the children attached to the scheduler. The parameters that control within-cir bandwidth allocation are the port-parent commands cir-level and cir-weight keywords. The cir-level keyword defines the port priority level that the scheduler or queue uses to receive bandwidth for its within-cir offered load. The cir-weight is used when multiple queues or schedulers exist at the same port priority level for within-cir bandwidth. The weight value defines the relative ratio that is used to distribute bandwidth at the priority level when more within-cir offered load exists than the port priority level has bandwidth.

A cir-weight equal to zero (the default value) has special meaning and informs the system that the queue or scheduler does not receive bandwidth from the within-cir distribution. Instead all bandwidth for the queue or scheduler must be allocated in the port scheduler’s above-cir pass.

Above-CIR Priority Level Parameters

The above-cir parameters define which port priority level the queue or scheduler should be associated with when receiving bandwidth for the queue’s or scheduler’s above-cir offered load. The above-cir offered load is the amount of bandwidth the queue or scheduler could use that is equal to or less than its defined PIR value (based on the queue or schedulers rate command) less any bandwidth that was given to the queue or scheduler during the above-cir scheduler pass. The parameters that control above-cir bandwidth allocation are the port-parent commands level and weight keywords. The level keyword defines the port priority level that the scheduler or queue uses to receive bandwidth for its above-cir offered load. The weight is used when multiple queues or schedulers exist at the same port priority level for above-cir bandwidth. The weight value defines the relative ratio that is used to distribute bandwidth at the priority level when more above-cir offered load exists than the port priority level has bandwidth.

CLI Syntax:
```
config>qos# scheduler-policy scheduler-policy-name
 tier {1 | 2 | 3}
    scheduler scheduler-name
       port-parent [level priority-level] [weight priority-
          weight] [cir-level cir-priority-level] [cir-weight
          cir-priority-weight]
```

CLI Syntax:
```
config>qos#
 sap-egress sap-egress-policy-id [create]
    queue queue-id [{auto-expedite | best-effort | expedite}]
       [priority-mode | profile-mode] [create]
       port-parent [level priority-level] [weight priority-
          weight] [cir-level cir-priority-level] [cir-weight
          cir-priority-weight]
```

CLI Syntax:
```
config>qos#
 network-queue network-queue-policy-name [create]
 no network-queue network-queue-policy-name
    queue queue-id [multipoint] [{auto-expedite | best-effort
       | expedite}] [priority-mode | profile-mode] [create]
       port-parent [level priority-level] [weight priority-
          weight] [cir-level cir-priority-level] [cir-weight
          cir-priority-weight]
```

Configuring Distributed LAG Rate

The following output displays a sample configuration and explanation with and without **dist-lag-rate-shared**.

```
*B:ALU-A>config>port# info
----------------------------------------------
        ethernet
            mode access
            egress-scheduler-policy "psp"
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------
*B:ALU-A>config>port# /configure lag 30
*B:ALU-A>config>lag# info
----------------------------------------------
        description "Description For LAG Number 30"
        mode access
        port 2/1/6
        port 2/1/10
        port 3/2/1
        port 3/2/2
        no shutdown

*B:ALU-A>config>service>ies>if>sap# /configure qos port-scheduler-policy "psp"
*B:ALU-A>config>qos>port-sched-plcy# info
----------------------------------------------
            max-rate 413202
```

Before enabling **dist-lag-rate-shared**, in the **port-scheduler-policy psp**, the max-rate achieved is twice 413202 kbps 816Mbps. This is because LAG has members from two different cards.

Two port-scheduler-instances are created, one on each card with the max-rate of 413202 kbps. This can be confirmed using the following show o/p.

Once dist-lag-rate-shared is enabled in port-scheduler-policy, this max-rate is enforced across all members of the LAG.

```
*B:ALU-A>config>service>ies>if>sap# /show qos scheduler-hierarchy sap lag-30 egress detail
===============================================================================
Scheduler Hierarchy - Sap lag-30
===============================================================================
Egress Scheduler Policy :
-------------------------------------------------------------------------------
Legend :
(*) real-time dynamic value
(w) Wire rates
B   Bytes
-------------------------------------------------------------------------------

Root (Egr)
| slot(2)
```

```
|--(S) : Tier0Egress:1->lag-30:0.0->1 (Port lag-30 Orphan)
|    |    AdminPIR:2000000    AdminCIR:0(sum)
|    |    Parent Limit Unused Bandwidth: not-found
|    |
|    |    AvgFrmOv:101.65(*)
|    |    AdminPIR:2000000(w) AdminCIR:0(w)
|    |
|    |    [Within CIR Level 0 Weight 0]
|    |    Assigned:0(w)       Offered:0(w)
|    |    Consumed:0(w)
|    |
|    |    [Above CIR Level 1 Weight 0]
|    |    Assigned:413202(w)  Offered:2000000(w) <----without dist-lag-rate-shared 413MB is
assigned to slot 2
|    |    Consumed:413202(w)
|    |
|    |
|    |    TotalConsumed:413202(w)
|    |    OperPIR:406494
|    |
|    |    [As Parent]
|    |    OperPIR:406494      OperCIR:0
|    |    ConsumedByChildren:406494
|    |
|    |
|    |--(Q) : 1->lag-30(2/1/6)->1
|    |    |    AdminPIR:1000000    AdminCIR:0
|    |    |    Parent Limit Unused Bandwidth: not-found
|    |    |    AvgFrmOv:101.65(*)
|    |    |    CBS:0 B             MBS:1310720 B
|    |    |    Depth:1045760 B     HiPrio:262144 B
|    |    |
|    |    |    [CIR]
|    |    |    Assigned:0          Offered:0
|    |    |    Consumed:0
|    |    |
|    |    |    [PIR]
|    |    |    Assigned:203247     Offered:1000000
|    |    |    Consumed:203247
|    |    |
|    |    |    OperPIR:205000      OperCIR:0
|    |    |
|    |    |    PktByteOffset:add 0*
|    |    |    OnTheWireRates:false
|    |    |    ATMOnTheWireRates:false
|    |    |    LastMileOnTheWireRates:false
|    |
|    |--(Q) : 1->lag-30(2/1/10)->1
|    |    |    AdminPIR:1000000    AdminCIR:0
|    |    |    Parent Limit Unused Bandwidth: not-found
|    |    |    AvgFrmOv:101.65(*)
|    |    |    CBS:0 B             MBS:1310720 B
|    |    |    Depth:1048320 B     HiPrio:262144 B
|    |    |
|    |    |    [CIR]
|    |    |    Assigned:0          Offered:0
|    |    |    Consumed:0
|    |    |
|    |    |    [PIR]
```

```
|   |   |    Assigned:203247     Offered:1000000
|   |   |    Consumed:203247
|   |   |
|   |   |    OperPIR:205000      OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
| slot(3)
|--(S) : Tier0Egress:1->lag-30:0.0->1 (Port lag-30 Orphan)
|   |    AdminPIR:2000000    AdminCIR:0(sum)
|   |    Parent Limit Unused Bandwidth: not-found
|   |
|   |    AvgFrmOv:101.65(*)
|   |    AdminPIR:2000000(w) AdminCIR:0(w)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0(w)       Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |    [Above CIR Level 1 Weight 0]
|   |    Assigned:413202(w)  Offered:2000000(w) <----without dist-lag-rate-shared 413MB is
assigned to slot 3
|   |    Consumed:413202(w)
|   |
|   |
|   |    TotalConsumed:413202(w)
|   |    OperPIR:406494
|   |
|   |    [As Parent]
|   |    OperPIR:406494      OperCIR:0
|   |    ConsumedByChildren:406494
|   |
|   |
|   |--(Q) : 1->lag-30(3/2/2)->1
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
|   |   |    CBS:0 B             MBS:1253376 B
|   |   |    Depth:1106976 B     HiPrio:147456 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:203247     Offered:1000000
|   |   |    Consumed:203247
|   |   |
|   |   |    OperPIR:203125      OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
|   |--(Q) : 1->lag-30(3/2/1)->1
```

```
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
|   |   |    CBS:0 B             MBS:1253376 B
|   |   |    Depth:1106976 B     HiPrio:147456 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:203247     Offered:1000000
|   |   |    Consumed:203247
|   |   |
|   |   |    OperPIR:203125      OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
```

The following output shows **dist-lag-rate-shared** enabled.

```
*B:ALU-A>config>qos>port-sched-plcy# dist-lag-rate-shared
*B:ALU-A>config>qos>port-sched-plcy# info
----------------------------------------------
            dist-lag-rate-shared
            max-rate 413202
----------------------------------------------
*B:ALU-A>config>qos>port-sched-plcy# !/show

*B:ALU-A>config>qos>port-sched-plcy# /show qos scheduler-hierarchy sap lag-30 egress detail
===============================================================================
Scheduler Hierarchy - Sap lag-30
===============================================================================
Egress Scheduler Policy :
-------------------------------------------------------------------------------
Legend :
(*) real-time dynamic value
(w) Wire rates
B   Bytes
-------------------------------------------------------------------------------
Root (Egr)
| slot(2)
|--(S) : Tier0Egress:1->lag-30:0.0->1 (Port lag-30 Orphan)
|   |    AdminPIR:2000000    AdminCIR:0(sum)
|   |    Parent Limit Unused Bandwidth: not-found
|   |
|   |    AvgFrmOv:101.65(*)
|   |    AdminPIR:2000000(w) AdminCIR:0(w)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0(w)       Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |    [Above CIR Level 1 Weight 0]
|   |    Assigned:206601(w)  Offered:2000000(w) <----with dist-lag-rate-shared 206 Mb is
assigned to slot 2
```

```
|   |    Consumed:206601(w)
|   |
|   |
|   |    TotalConsumed:206601(w)
|   |    OperPIR:203247
|   |
|   |    [As Parent]
|   |    OperPIR:203247      OperCIR:0
|   |    ConsumedByChildren:203247
|   |
|   |
|   |--(Q) : 1->lag-30(2/1/6)->1
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
|   |   |    CBS:0 B             MBS:1310720 B
|   |   |    Depth:1045760 B     HiPrio:262144 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:101624     Offered:1000000
|   |   |    Consumed:101624
|   |   |
|   |   |    OperPIR:102000      OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
|   |--(Q) : 1->lag-30(2/1/10)->1
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
|   |   |    CBS:0 B             MBS:1310720 B
|   |   |    Depth:1047040 B     HiPrio:262144 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:101624     Offered:1000000
|   |   |    Consumed:101624
|   |   |
|   |   |    OperPIR:102000      OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
| slot(3)
|--(S) : Tier0Egress:1->lag-30:0.0->1 (Port lag-30 Orphan)
|   |    AdminPIR:2000000    AdminCIR:0(sum)
|   |    Parent Limit Unused Bandwidth: not-found
```

```
|   |
|   |    AvgFrmOv:101.65(*)
|   |    AdminPIR:2000000(w) AdminCIR:0(w)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0(w)       Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |    [Above CIR Level 1 Weight 0]
|   |    Assigned:206601(w)  Offered:2000000(w) <----with dist-lag-rate-shared 206 Mb is
assigned to slot 3
|   |    Consumed:206601(w)
|   |
|   |
|   |    TotalConsumed:206601(w)
|   |    OperPIR:203247
|   |
|   |    [As Parent]
|   |    OperPIR:203247      OperCIR:0
|   |    ConsumedByChildren:203247
|   |
|   |
|   |--(Q) : 1->lag-30(3/2/2)->1
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
|   |   |    CBS:0 B             MBS:1253376 B
|   |   |    Depth:1105728 B     HiPrio:147456 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:101624     Offered:1000000
|   |   |    Consumed:101624
|   |   |
|   |   |    OperPIR:101500      OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
|   |--(Q) : 1->lag-30(3/2/1)->1
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
|   |   |    CBS:0 B             MBS:1253376 B
|   |   |    Depth:1105728 B     HiPrio:147456 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:101624     Offered:1000000
|   |   |    Consumed:101624
|   |   |
```

```
|   |   |    OperPIR:101500      OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
===============================================================================
*B:ALU-A>config>qos>port-sched-plcy#
```

If one of the member links of the LAG goes down, then the max-rate is divided among the remaining lag members.

Card 2 is assigned 137734 (1/3 of max-rate 413202)

Card 3 is assigned 275468 (2/3 of max-rate 413202)

```
*B:ALU-A>config>qos>port-sched-plcy# /show lag 30 detail
===============================================================================
LAG Details
===============================================================================
Description        : Description For LAG Number 30
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Lag-id              : 30                    Mode                 : access
Adm                 : up                    Opr                  : up
Thres. Exceeded Cnt : 11                    Port Threshold       : 0
Thres. Last Cleared : 06/25/2014 21:47:49   Threshold Action     : down
Dynamic Cost        : false                 Encap Type           : null
Configured Address  : 00:1a:f0:1d:8b:c9     Lag-IfIndex          : 1342177310
Hardware Address    : 00:1a:f0:1d:8b:c9     Adapt Qos (access)   : distribute
Hold-time Down      : 0.0 sec               Port Type            : standard
Per-Link-Hash       : disabled
Include-Egr-Hash-Cfg: disabled
Per FP Ing Queuing  : disabled              Per FP Egr Queuing   : disabled
Per FP SAP Instance : disabled
LACP                : disabled
Standby Signaling   : lacp
Port weight         : 0 gbps                Number/Weight Up     : 3
Weight Threshold    : 0                     Threshold Action     : down

-------------------------------------------------------------------------------
Port-id        Adm     Act/Stdby Opr     Primary   Sub-group     Forced  Prio
-------------------------------------------------------------------------------
2/1/6          up      active    down    yes       1            -       32768
2/1/10         up      active    up                1            -       32768
3/2/1          up      active    up                1            -       32768
3/2/2          up      active    up                1            -       32768
===============================================================================
*B:ALU-A>config>qos>port-sched-plcy# /show qos scheduler-hierarchy sap lag-30 egress detail

===============================================================================
Scheduler Hierarchy - Sap lag-30
===============================================================================
Egress Scheduler Policy :
```

```
-------------------------------------------------------------------------------
Legend :
(*) real-time dynamic value
(w) Wire rates
B   Bytes
-------------------------------------------------------------------------------

Root (Egr)
| slot(2)
|--(S) : Tier0Egress:1->lag-30:0.0->1 (Port lag-30 Orphan)
|   |    AdminPIR:1000000    AdminCIR:0(sum)
|   |    Parent Limit Unused Bandwidth: not-found
|   |
|   |    AvgFrmOv:101.65(*)
|   |    AdminPIR:1000000(w) AdminCIR:0(w)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0(w)       Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |    [Above CIR Level 1 Weight 0]
|   |    Assigned:137734(w)  Offered:1000000(w)
|   |    Consumed:137734(w)
|   |
|   |
|   |    TotalConsumed:137734(w)
|   |    OperPIR:135498
|   |
|   |    [As Parent]
|   |    OperPIR:135498      OperCIR:0
|   |    ConsumedByChildren:135498
|   |
|   |
|   |--(Q) : 1->lag-30(2/1/6)->1
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
|   |   |    CBS:0 B             MBS:1310720 B
|   |   |    Depth:0 B           HiPrio:262144 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:67749      Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    OperPIR:68000       OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
|   |--(Q) : 1->lag-30(2/1/10)->1
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
```

```
|   |   |    CBS:0 B             MBS:1310720 B
|   |   |    Depth:1044480 B     HiPrio:262144 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:135498     Offered:1000000
|   |   |    Consumed:135498
|   |   |
|   |   |    OperPIR:135000      OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
| slot(3)
|--(S) : Tier0Egress:1->lag-30:0.0->1 (Port lag-30 Orphan)
|   |    AdminPIR:2000000    AdminCIR:0(sum)
|   |    Parent Limit Unused Bandwidth: not-found
|   |
|   |    AvgFrmOv:101.65(*)
|   |    AdminPIR:2000000(w) AdminCIR:0(w)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0(w)       Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |    [Above CIR Level 1 Weight 0]
|   |    Assigned:275468(w)  Offered:2000000(w)
|   |    Consumed:275468(w)
|   |
|   |
|   |    TotalConsumed:275468(w)
|   |    OperPIR:270996
|   |
|   |    [As Parent]
|   |    OperPIR:270996      OperCIR:0
|   |    ConsumedByChildren:270996
|   |
|   |
|   |--(Q) : 1->lag-30(3/2/2)->1
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
|   |   |    CBS:0 B             MBS:1253376 B
|   |   |    Depth:1106976 B     HiPrio:147456 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:135498     Offered:1000000
|   |   |    Consumed:135498
|   |   |
|   |   |    OperPIR:135625      OperCIR:0
```

```
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |
|   |--(Q) : 1->lag-30(3/2/1)->1
|   |   |    AdminPIR:1000000    AdminCIR:0
|   |   |    Parent Limit Unused Bandwidth: not-found
|   |   |    AvgFrmOv:101.65(*)
|   |   |    CBS:0 B             MBS:1253376 B
|   |   |    Depth:1105728 B     HiPrio:147456 B
|   |   |
|   |   |    [CIR]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [PIR]
|   |   |    Assigned:135498     Offered:1000000
|   |   |    Consumed:135498
|   |   |
|   |   |    OperPIR:135625      OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
|   |

===============================================================================
*B:ALU-A>config>qos>port-sched-plcy#
```

The following output shows the **max-rate percent** value.

```
*B:ALU-A>config>qos>port-sched-plcy# info
----------------------------------------------
            max-rate percent 50.00
----------------------------------------------
```

With **max-rate percent**, the **max-rate** is capped to the percent of the active LAG capacity.

When **max-rate** is configured as percentage and the **dist-lag-rate-shared** is ignored.

The group rate, level pir and cir rate can be entered as percent.

```
*B:ALU-A>config>qos>port-sched-plcy# info
----------------------------------------------
 dist-lag-rate-shared
  max-rate percent 30.00
 group "test" create
                percent-rate 20.00 cir 20.00
  exit
  level 1 percent-rate 10.00 percent-cir 10.00
  level 2 percent-rate 20.00 percent-cir 20.00
  level 3 percent-rate 30.00 percent-cir 30.00
```

```
  level 4 percent-rate 40.00 percent-cir 40.00
  level 5 percent-rate 50.00 percent-cir 50.00
level 6 percent-rate 60.00 percent-cir 60.00
  level 7 percent-rate 70.00 percent-cir 70.00
  level 8 percent-rate 80.00 percent-cir 80.00
```

Port scheduler-Overrides

Both max-rate and level can be overridden if they are of the same type as in the policy being overridden.

```
*B:ALU-A>config>port>ethernet>egr-sched-override$ info
----------------------------------------------
                max-rate percent 50.00
                level 1 percent-rate 10.00 percent-cir 10.00
                level 2 percent-rate 20.00 percent-cir 20.00
                level 3 percent-rate 30.00 percent-cir 30.00
                level 4 percent-rate 40.00 percent-cir 40.00
                level 5 percent-rate 50.00 percent-cir 50.00
                level 6 percent-rate 60.00 percent-cir 60.00
                level 7 percent-rate 70.00 percent-cir 70.00
                level 8 percent-rate 80.00 percent-cir 80.00
```

The following are additions to the to the show command output:

Dist Lag Rate, Lvl and Group PIR and Cir Percent rates

```
*B:ALU-A>config>qos>port-sched-plcy# /show qos port-scheduler-policy "psp2"
===============================================================================
QoS Port Scheduler Policy
===============================================================================
Policy-Name       : psp2
Description       : (Not Specified)
Max Rate          : max                Max Rate Percent  : 30.00
Dist LAG Rate     : True               Last changed      : 07/16/2014 21:31:51
Group             : test
Group PIR         : max                Group CIR         : max
Group PIR Percent : 20.00              Group CIR Percent : 20.00

Lvl1 PIR          : max                Lvl1 CIR          : max
Lvl1 PIR Percent  : 10.00              Lvl1 CIR Percent  : 10.00
Lvl2 PIR          : max                Lvl2 CIR          : max
Lvl2 PIR Percent  : 20.00              Lvl2 CIR Percent  : 20.00
Lvl3 PIR          : max                Lvl3 CIR          : max
Lvl3 PIR Percent  : 30.00              Lvl3 CIR Percent  : 30.00
Lvl4 PIR          : max                Lvl4 CIR          : max
Lvl4 PIR Percent  : 40.00              Lvl4 CIR Percent  : 40.00
Lvl5 PIR          : max                Lvl5 CIR          : max
Lvl5 PIR Percent  : 50.00              Lvl5 CIR Percent  : 50.00
Lvl6 PIR          : max                Lvl6 CIR          : max
Lvl6 PIR Percent  : 60.00              Lvl6 CIR Percent  : 60.00
Lvl7 PIR          : max                Lvl7 CIR          : max
Lvl7 PIR Percent  : 70.00              Lvl7 CIR Percent  : 70.00
Lvl8 PIR          : max                Lvl8 CIR          : max
Lvl8 PIR Percent  : 80.00              Lvl8 CIR Percent  : 80.00
```

```
Orphan Lvl        : default             Orphan Weight     : default
Orphan CIR-Lvl    : default             Orphan CIR-Weight : default

--snip--
-------------------------------------------------------------------------------
Egr Port Sched Override
-------------------------------------------------------------------------------
Max Rate          : max*                      Max Rate Percent : 50.00
Lvl1 PIR          : max*                      Lvl1 CIR         : max*
Lvl1 PIR Percent  : 10.00                     Lvl1 CIR Percent : 10.00
Lvl2 PIR          : max*                      Lvl2 CIR         : max*
Lvl2 PIR Percent  : 20.00                     Lvl2 CIR Percent : 20.00
Lvl3 PIR          : max*                      Lvl3 CIR         : max*
Lvl3 PIR Percent  : 30.00                     Lvl3 CIR Percent : 30.00
Lvl4 PIR          : max*                      Lvl4 CIR         : max*
Lvl4 PIR Percent  : 40.00                     Lvl4 CIR Percent : 40.00
Lvl5 PIR          : max*                      Lvl5 CIR         : max*
Lvl5 PIR Percent  : 50.00                     Lvl5 CIR Percent : 50.00
Lvl6 PIR          : max*                      Lvl6 CIR         : max*
Lvl6 PIR Percent  : 60.00                     Lvl6 CIR Percent : 60.00
Lvl7 PIR          : max*                      Lvl7 CIR         : max*
Lvl7 PIR Percent  : 70.00                     Lvl7 CIR Percent : 70.00
Lvl8 PIR          : max*                      Lvl8 CIR         : max*
Lvl8 PIR Percent  : 80.00                     Lvl8 CIR Percent : 80.00
* means the value is inherited
-------------------------------------------------------------------------------
```

Service Management Tasks

This section discusses the following service management tasks:

- Deleting QoS Policies on page 545
- Copying and Overwriting Scheduler Policies on page 548
- Editing QoS Policies on page 550

Deleting QoS Policies

There are no scheduler or port-scheduler policies associated with customer or service entities. Removing a scheduler or port-scheduler policy from a multi-service customer site causes the created schedulers to be removed which makes them unavailable for the ingress SAP queues associated with the customer site. Queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have ingress queues that rely on the schedulers enter into an orphaned state on one or more queues.

A QoS scheduler policy cannot be deleted until it is removed from all customer multi-service sites or service SAPs where it is applied.

```
SR7>config>qos# no scheduler-policy SLA2
MINOR: QOS #1003 The policy has references
SR7>config>qos#
```

Removing a QoS Policy from a Customer Multi-Service Site

CLI Syntax:
```
config>service>customer customer-id
    multi-service-site customer-site-name
        egress
            no scheduler-policy
        ingress
            no scheduler-policy
```

Example:
```
config>service>customer# multi-service-site "Test"
config>service>cust>multi-service-site# ingress
config>service>cust>multi-service-site>ingress# no
      scheduler-policy
```

Removing a QoS Policy from SAP(s)

CLI Syntax:
```
config>service# {epipe|vpls} service-id [customer customer-id]
 sap sap-id
    egress
       no scheduler policy
    ingress
       no scheduler policy
```

Example:
```
config>service# epipe 6
config>service>epipe# sap sap 1/1/9:0
config>service>epipe>sap# egress
config>service>epipe>sap>egress# no scheduler-policy
config>service>epipe>sap>egress# exit
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress#
config>service>epipe>sap>ingress# no scheduler-policy
```

CLI Syntax:
```
config>service# {ies|vprn} service-id [customer customer-id]
 interface ip-int-name
    sap sap-id
       egress
          no scheduler policy
       ingress
          no scheduler policy
```

Example:
```
config>service# vprn 1
onfig>service>vprn# interface "to-ce1"
config>service>vprn>if# sap 1/1/10:1
config>service>vprn>if>sap# ingress
config>service>vprn>if>sap>ingress# no scheduler-policy
config>service>vprn>if>sap>ingress# exit
config>service>vprn>if>sap# egress
config>service>vprn>if>sap>egress# no scheduler-policy
config>service>vprn>if>sap>egress# exit
config>service>vprn>if>sap#
```

Removing a Policy from the QoS Configuration

To delete a scheduler policy, enter the following commands:

CLI Syntax: `config>qos# no scheduler-policy` *`network-policy-id`*

Example: `config>qos# no scheduler-policy` ***`SLA1`***

To delete a port scheduler policy, enter the following commands:

CLI Syntax: `config>qos# no port-scheduler-policy` *`network-policy-id`*

Example: `config>qos# no port-scheduler-policy` ***`test1`***

Copying and Overwriting Scheduler Policies

You can copy an existing QoS policy, rename it with a new QoS policy value, or overwrite an existing policy. The `overwrite` option must be specified or an error occurs if the destination policy exists.

CLI Syntax: `config>qos> copy scheduler-policy` *`src-name dst-name`* `[overwrite]`

Example: `config>qos# copy scheduler-policy SLA1 SLA2`

```
A:SR>config>qos#
...
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
        scheduler-policy "SLA1" create
            description "NetworkControl(3), Voice(2) and NonVoice(1) have strict priorities"
            tier 1
                scheduler "All_traffic" create
                    description "All traffic goes to this scheduler eventually"
                    rate 11000
                exit
            exit
            tier 2
                scheduler "NetworkControl" create
                    description "network control traffic within the VPN"
                    parent "All_traffic" level 3 cir-level 3
                    rate 100
                exit
                scheduler "NonVoice" create
                    description "NonVoice of VPN and Internet traffic will be serviced by
this scheduler"
                    parent "All_traffic" cir-level 1
                    rate 11000
                exit
                scheduler "Voice" create
                   description "Any voice traffic from VPN and Internet use this scheduler"
                    parent "All_traffic" level 2 cir-level 2
                    rate 5500
                exit
            exit
            tier 3
                scheduler "Internet_be" create
                    parent "NonVoice" cir-level 1
                exit
                scheduler "Internet_priority" create
                    parent "NonVoice" level 2 cir-level 2
                exit
...
        scheduler-policy "SLA2" create
            description "NetworkControl(3), Voice(2) and NonVoice(1) have strict priorities"
            tier 1
                scheduler "All_traffic" create
                    description "All traffic goes to this scheduler eventually"
```

```
                    rate 11000
                exit
            exit
            tier 2
                scheduler "NetworkControl" create
                    description "network control traffic within the VPN"
                    parent "All_traffic" level 3 cir-level 3
                    rate 100
                exit
                scheduler "NonVoice" create
                    description "NonVoice of VPN and Internet traffic will be serviced by
this scheduler"
                    parent "All_traffic" cir-level 1
                    rate 11000
                exit
                scheduler "Voice" create
                   description "Any voice traffic from VPN and Internet use this scheduler"
                    parent "All_traffic" level 2 cir-level 2
                    rate 5500
                exit
            exit
            tier 3
                scheduler "Internet_be" create
                    parent "NonVoice" cir-level 1
                exit
                scheduler "Internet_priority" create
                    parent "NonVoice" level 2 cir-level 2
                exit
...
#---------------------------------------
A:SR>config>qos#
```

Editing QoS Policies

You can change existing policies and entries in the CLI. The changes are applied immediately to all customer multi-service sites and service SAPs where the policy is applied. To prevent configuration errors use the copy command to make a duplicate of the original policy to a work area, make the edits, and then overwrite the original policy.

QoS Scheduler Policy Command Reference

Command Hierarchies

- Scheduler Policy Configuration Commands on page 551
- Port Scheduler Policy Configuration Commands on page 551
- Operational Commands on page 552
- Show Commands on page 552
- Clear Commands on page 553

Scheduler Policy Configuration Commands

```
config
    — qos
        — [no] scheduler-policy scheduler-policy-name
            — description description-string
            — no description
            — [no] frame-based-accounting
            — parent-location {none | sub | vport}
            — no parent-location
            — [no] tier tier
                — no] scheduler scheduler-name
                    — description description-string
                    — no description
                    — [no] limit-unused-bandwidth
                    — parent scheduler-name [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]
                    — no parent
                    — port-parent [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]
                    — no port-parent
                    — rate [pir-rate] [cir cir-rate]
                    — no rate
```

Port Scheduler Policy Configuration Commands

```
config
    — qos
        — [no] port-scheduler-policy port-scheduler-name
            — description description-string
            — no description
            — [no] dist-lag-rate-shared
            — group name [create]
```

- **no group** *name*
  - **percent-rate** *pir-percent* [**cir** *cir-percent*]
  - **no percent-rate**
  - **rate** *pir-rate* [**cir** *cir-rate*]
  - **no rate**
- **level** *priority-level* **rate** *pir-rate* [**cir** *cir-rate*] **group** *name* [**weight** *weight*]
- **level** *priority-level* **percent-rate** *pir-percent* [**percent-cir** *cir-percent*] **group** *name* [**weight** *weight*]
- **level** *priority-level* **rate** *pir-rate* [**cir** *cir-rate*]
- **level** *priority-level* **percent-rate** *pir-percent* [**percent-cir** *cir-percent*]
- **no level** *priority-level*
- **level** *priority-level* **rate** *pir-rate* [**cir** *cir-rate*]
- **no level** *priority-level*
- **max-rate** *pir-rate*
- **max-rate** **percent** *percent-rate*
- **no max-rate**
- **orphan-override** [**level** *priority-level*] [**weight** *percent*] [**cir-level** *priority-level*] [**cir-weight** *cir-weight*]
- **no orphan-override**

Operational Commands

config
- **qos**
  - **copy** **scheduler-policy** *src-name dst-name* [**overwrite**]
  - **copy** **port-scheduler-policy** *src-name dst-name* [**overwrite**]

Show Commands

show
- **qos**
  - **scheduler-hierarchy customer** *customer-id* **site** *customer-site-name* [**scheduler** *scheduler-name*] [**ingress** | **egress**] [**detail**]
  - **scheduler-hierarchy port** *port-id* [**detail**]
  - **scheduler-hierarchy port** *port-id* [**detail**] **queue-group** *queue-group-name* [**scheduler** *scheduler-name*] [**ingress** | **egress**]
  - **scheduler-hierarchy sap** *sap-id* [**scheduler** *scheduler-name*] [**ingress** | **egress**] [**detail**]
  - **scheduler-hierarchy subscriber** *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress** | **egress**] [**detail**]
  - **scheduler-name** *scheduler-name*
  - **scheduler-policy** *scheduler-name* [**association** | **sap-ingress** *policy-id* | **sap-egress** *policy-id*]
  - **scheduler-stats customer** *customer-id* **site** *customer-site-name* [**scheduler** *scheduler-name*] [**ingress**|**egress**]
  - **scheduler-stats sap** *sap-id* [**scheduler** *scheduler-name*] [**ingress**|**egress**]
  - **scheduler-stats subscriber** *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress**|**egress**]

show
- **qos**
  - **port-scheduler-policy** [*port-scheduler-policy-name*] [**association**]

— **port-scheduler-policy** *port-scheduler-policy-name* **network-policy** *network-queue-policy-name*
— **port-scheduler-policy** *port-scheduler-policy-name* **sap-egress** *policy-id*
— **port-scheduler-policy** *port-scheduler-policy-name* **scheduler-policy** *scheduler-policy-name*
— **port-scheduler-policy** *port-scheduler-policy-name* **scheduler-policy** *scheduler-policy-name* **sap-egress** *policy-id*

Clear Commands

clear
— **qos**
— **scheduler-stats**
— **sap** *sap-id* [**scheduler** *scheduler-name*] [**ingress** | **egress**]

Generic Commands

description

Syntax **description** *description-string*
no description

Context config>qos>scheduler-policy
config>qos>scheduler-policy>tier>scheduler
config>qos>port-scheduler-policy

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax **copy scheduler-policy** *src-name dst-name* [**overwrite**]
copy port-scheduler-policy *src-name dst-name* [**overwrite**]

Context config>qos

Description This command copies existing QoS policy entries for a QoS policy to another QoS policy.

The **copy** command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

If **overwrite** is not specified, an error will occur if the destination policy exists.

Parameters **scheduler-policy** *src-name dst-name* — Indicates that the source policy and the destination policy are scheduler policy. Specify the source policy that the copy command will attempt to copy from and specify the destination policy to which the command will copy a duplicate of the policy.

port-scheduler-policy *src-name dst-name* — Indicates that the source policy and the destination policy are port scheduler policy IDs. Specify the source policy that the copy command will attempt to copy from and specify the destination policy name to which the command will copy a duplicate of the policy.

overwrite — Forces the destination policy name to be copied as specified. When forced, everything in the existing destination policy will be completely overwritten with the contents of the source policy.

Scheduler Policy Commands

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy *scheduler-policy-name*

Context config>qos

Description Each scheduler policy is divided up into groups of schedulers based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations.

The **scheduler-policy** command creates a scheduler policy or allows you to edit an existing policy.The policy defines the hierarchy and operating parameters for virtual schedulers. Merely creating a policy does not create the schedulers; it only provides a template for the schedulers to be created when the policy is associated with a SAP or multi-service site.

Each scheduler policy must have a unique name within the context of the system. Modifications made to an existing policy are executed on all schedulers that use the policy. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If a **scheduler-policy-name** does not exist, it is assumed that an attempt is being made to create a new policy. The success of the command execution is dependent on the following:

1. The maximum number of scheduler policies has not been configured.
2. The provided scheduler-policy-name is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of scheduler policies has been exceeded a configuration error occurs, the command will not execute, and the CLI context will not change.

If the provided scheduler-policy-name is invalid according to the criteria below, a name syntax error occurs, the command will not execute, and the CLI context will not change.

Default **none** — Each scheduler policy must be explicitly created.

Parameters *scheduler-policy-name* — The name of the scheduler policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

frame-based-accounting

Syntax **frame-based-accounting**
no frame-based-accounting

Context config>qos>scheduler-policy

Description The frame-based-accounting command is used to enable frame based for both the children queues parented to the scheduling policy and for the schedulers within the scheduler policy.

Once frame based accounting is enabled on the policy, all queues associated with the scheduler (through the parent command on each queue) will have their rate and CIR values interpreted as frame based values. When shaping, the queues will include the 12 byte Inter-Frame Gap (IFG) and 8 byte preamble for each packet scheduled out the queue. The profiling CIR threshold will also include the 20 byte frame encapsulation overhead. Statistics associated with the queue do not include the frame encapsulation overhead.

The scheduler policy's scheduler rate and CIR values will be interpreted as frame based values.

The configuration of *parent-location* and **frame-based-accounting** in a scheduler policy is mutually exclusive to ensure consistency between the different scheduling levels.

The **no** frame-based-accounting command is used to return all schedulers within the policy and queues associated with the policy to the default packet based accounting mode. If frame based accounting is not currently enabled for the scheduling policy, the no frame-based-accounting command has no effect.

parent-location

Syntax **parent-location** {**none** | **sub** | **vport**}
no parent-location

Context config>qos>scheduler-policy

Description This command determines the expected location of the parent schedulers for the tier 1 schedulers configured with a parent command within the scheduler-policy. The parent schedulers must be configured within a scheduler-policy applied at the location corresponding to the parent-location parameter.

If a parent scheduler name does not exist at the specified location, the schedulers will not be parented and will be orphaned.

The configuration of parent-location and frame-based accounting in a scheduler policy is mutually exclusive in order to ensure consistency between the different scheduling levels.

The **no** form of the command reverts to the default.

Default none

Parameters **none —** This parameter indicates that the tier 1 schedulers do not have a parent scheduler and the configuration of the parent under a tier 1 scheduler is blocked. Conversely, this parameter is blocked when any tier 1 scheduler has a parent configured.

sub — When the scheduler-policy is applied to an sla-profile for a subscriber, the parent schedulers of the tier 1 schedulers need to be configured in the scheduler-policy applied to the subscriber's sub-profile.

If this parameter is configured within a scheduler-policy that is applied to any object except for the egress of an sla-profile, the configured parent schedulers will not be found and so the tier 1 schedulers will not be parented and will be orphaned.

vport — When the scheduler-policy is applied to an sla-profile, a sub-profile for a subscriber or to the egress of a pseudowire SAP, the parent schedulers of the tier 1 schedulers need to be configured in the scheduler-policy applied to the vport to which the subscriber will be assigned.

If this parameter is configured within a scheduler-policy that is applied to to any object except for the egress of an sla-profile or sub-profile, or to the egress of a PW SAP, the configured parent schedulers will not be found and so the tier 1 schedulers will not be parented and will be orphaned.

tier

Syntax **tier** *tier*

Context config>qos>scheduler-policy

Description This command identifies the level of hierarchy that a group of schedulers are associated with. Within a tier level, a scheduler can be created or edited. Schedulers created within a tier can only be a child (take bandwidth from a scheduler in a higher tier). Tier levels increase sequentially with 1 being the highest tier. All tier 1 schedulers are considered to be root and cannot be a child of another scheduler. Schedulers defined in tiers other than 1 can also be root (parentless).

3 tiers (levels 1, 2 and 3) are supported.

The **save config** and **show config** commands only display information on scheduler tiers that contain defined schedulers. When all schedulers have been removed from a level, that level ceases to be included in output from these commands.

Parameters *tier* — This parameter is required to indicate the group of schedulers to create or be edited. Tier ***levels*** cannot be created or deleted. If a value for level is given that is out-of-range, an error will occur and the current context of the CLI session will not change.

Values 1 — 3

Default None

scheduler

Syntax **scheduler** *scheduler-name*
no scheduler *scheduler-name*

Context config>qos>scheduler-policy>tier ***level***

Description This command creates a new scheduler or edits an existing scheduler within the scheduler policy tier. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the

scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If ***scheduler-name*** already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the ***scheduler-name*** exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change.

If the ***scheduler-name*** does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

1. The maximum number of schedulers has not been configured.
2. The provided ***scheduler-name*** is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.

Parameters *scheduler-name* — The name of the scheduler.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Default None. Each scheduler must be explicitly created.

create — This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given ***scheduler-name***. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

limit-unused-bandwidth

Syntax **[no] limit-unused-bandwidth**

Context config>qos>scheduler-policy>tier>scheduler

Description This command is used to enable (or disable) aggregate rate overrun protection on the agg-rate context.

parent

Syntax **parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
no parent

Context config>qos>scheduler-policy>tier ***level***>scheduler ***scheduler-name***

Description This command defines an optional parent scheduler that is higher up the policy hierarchy. Only schedulers in tier levels 2 and 3 can have a parental association. When multiple schedulers and/or queues share a child status with the scheduler on the parent, the weight or strict parameters define how this scheduler contends with the other children for the parent's bandwidth. The parent scheduler can be removed or changed at anytime and is immediately reflected on the schedulers created by association of this scheduler policy.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of the command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child scheduler attempts to operate based on its configured rate parameter. Removing the parent association on the scheduler within the policy will take effect immediately on all schedulers with ***scheduler-name*** that have been created using the ***scheduler-policy-name***.

Parameters *scheduler-name* — The ***scheduler-name*** must already exist within the context of the scheduler policy in a tier that is higher (numerically lower).

Values Any valid **scheduler-name** existing on a higher tier within the scheduler policy.

Default None. Each parental association must be explicitly created.

weight *weight* **—** **Weight** defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict level defined by the **level** parameter. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.

A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

Default 1

level *level* **—** The **level** keyword defines the strict priority level of this scheduler in comparison to other child schedulers and queues vying for bandwidth on the parent scheduler-name during the 'above CIR' distribution phase of bandwidth allocation. During the above CIR distribution phase, any queues or schedulers defined at a lower strict level receive no parental bandwidth until all queues and schedulers defined with a higher (numerically larger) strict level on the parent have reached their maximum bandwidth or have satisfied their offered load requirements.

When the similar **cir-level** parameter default (undefined) are retained for the child scheduler, bandwidth is only allocated to the scheduler during the above CIR distribution phase.

Children of the parent scheduler with a lower strict priority level will not receive bandwidth until all children with a higher strict priority level have either reached their maximum bandwidth or are idle. Children with the same strict level are serviced according to their weight.

Values 1 — 8

Default 1

cir-weight *cir-weight* **—** The **cir-weight** keyword defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same ***cir-level*** defined by the **cir-level** parameter. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the queue or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.

A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 — 100

Default 1

cir-level *cir-level* **—** The **cir-level** keyword defines the strict priority CIR level of this scheduler in comparison to other child schedulers and queues vying for bandwidth on the parent ***scheduler-name*** during the 'within CIR' distribution phase of bandwidth allocation. During the 'within CIR' distribution phase, any queues or schedulers defined at a lower strict CIR level receive no parental bandwidth until all queues and schedulers defined with a higher (numerically larger) strict CIR level on the parent have reached their CIR bandwidth or have satisfied their offered load requirements.

If the scheduler's **cir-level** parameter retains the default (undefined) state, bandwidth is only allocated to the scheduler during the above CIR distribution phase.

Children with the same strict cir-level are serviced according to their cir-weight.

Values Undefined, 1 — 8

Default Undefined

port-parent

Syntax **port-parent** [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
no port-parent

Context config>qos>scheduler-policy>tier>scheduler

Description The **port-parent** command defines a child/parent association between an egress queue and a port based scheduler or between an intermediate service scheduler and a port based scheduler. The command may be issued in three distinct contexts; **sap-egress queue** *queue-id*, **network-queue queue** *queue-id* and **scheduler-policy scheduler** *scheduler-name*. The **port-parent** command allows for a set of within-CIR and above-CIR parameters that define the port priority levels and weights for the queue or scheduler. If the **port-parent** command is executed without any parameters, the default parameters are assumed.

In this context, the **port-parent** command is mutually exclusive to the **parent** command (used to create a parent/child association between a queue and an intermediate scheduler). Executing a **port-parent** command when a parent definition is in place causes the current intermediate scheduler association to be removed and replaced by the defined port-parent association. Executing a **parent** command when a port-parent definition exists causes the port scheduler association to be removed and replaced by the defined intermediate scheduler name.

Changing the parent context on a SAP egress policy queue may cause a SAP or subscriber context of the queue (policy associated with a SAP or subscriber profile) to enter an orphaned state. If an instance of a queue is created on a port or channel that does not have a port scheduler enabled and the sap-egress policy creating the queue has a port-parent association, the queue will be allowed to run according to its own rate parameters and will not be controlled by a virtual scheduling context. If an instance of a queue is on a port or channel that has a port scheduler configured and the sap-egress policy defines the queue as having a non-existent intermediate scheduler parent, the queue will be treated as an orphan and will be handled according to the current orphan behavior on the port scheduler.

The **no** form of this command removes a port scheduler parent association for the queue or scheduler. If a port scheduler is defined on the port which the queue or scheduler instance exists, the queue or scheduler will become orphaned if an port scheduler is configured on the egress port of the queue or scheduler.

Default **no port-parent**

Parameters **weight** *weight* **—** Defines the weight the queue or scheduler will use at the above-cir port priority level (defined by the level parameter).

Values 0 — 100

Default 1

level *level* **—** Defines the port priority the queue or scheduler will use to receive bandwidth for its above-cir offered-load.

Values 1 — 8 (8 is the highest priority)

Default 1

cir-weight *cir-weight* **—** Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value),

the queue or scheduler does not receive bandwidth during the port scheduler's within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 100

cir-level *cir-level* **—** Defines the port priority the queue or scheduler will use to receive bandwidth for its within-cir offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port scheduler's within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 — 8 (8 is the highest priority)

Default 0

rate

Syntax **rate** [*pir-rate*] [**cir** *cir-rate*]
no rate

Context config>qos>scheduler-policy>tier>scheduler

Description The **rate** command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other then its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns all queues created with this *queue-id* by association with the QoS policy to the default PIR and CIR parameters.

Parameters *pir-rate* — The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue will operate. A value of 0 to 100000000 or the keyword **max** is accepted. Any other value will result in an error without modifying the current PIR rate.

To calculate the actual PIR rate, the rate described by the queue's **rate** is multiplied by the *pir-rate*.

The SAP ingress context for PIR is independent of the defined forwarding class (fc) for the queue. The default **pir** and definable range is identical for each class. The PIR in effect for a queue defines the maximum rate ate which the queue will be allowed to forward packets in a given second, thus shaping the queue's output.

The PIR parameter for SAP ingress queues do not have a negate (**no**) function. To return the queue's PIR rate to the default value, that value must be specified as the PIR value.

Values 1 — 100000000, **max**

Default max

cir cir-rate — The **cir** parameter accepts a step-multiplier value that specifies the multiplier used to determine the CIR rate at which the queue will operate. A value of 0 to 250 or the keyword **max** is accepted. Any other value will result in an error without modifying the current CIR rate.

To calculate the actual CIR rate, the rate described by the **rate pir** *pir-rate*is multiplied by the cir *cir-rate*. If the **cir** is set to max, then the CIR rate is set to infinity.

The SAP ingress context for CIR is dependent on the defined forwarding class (fc) for the queue. The default CIR and definable range is different for each class. The CIR in effect for a queue defines both its profile (in or out) marking level as well as the relative importance compared to other queues for scheduling purposes during congestion periods.

Values 0 — 100000000, **max**, **sum**

Default sum

Port Scheduler Policy Commands

port-scheduler-policy

Syntax [**no**] **port-scheduler-policy** *port-scheduler-name*

Context config>qos

Description When a port scheduler has been associated with an egress port, it is possible to override the following parameters:

- The max-rate allowed for the scheduler
- The maximum rate for each priority level (8 through 1)
- The cir associated with each priority level (8 through 1)

The orphan priority level (level 0) has no configuration parameters and cannot be overridden.

The **no** form of the command removes a port scheduler policy from the system. If the port scheduler policy is associated with an egress port or channel, the command will fail.

Parameters *port-scheduler-name* — Specifies an existing port scheduler name. Each port scheduler must be uniquely named within the system and can be up to 32 ASCII characters in length.

dist-lag-rate-shared

Syntax [**no**] **dist-lag-rate-shared**

Context config>qos>port-scheduler-policy

Description This command enables sharing of rates when the port on which this port-scheduler-policy is configured is part of a LAG configured in **distribute** mode.

When enabled, the absolute rate values configured as part of the max-rate, PIR/CIR group rates and PIR/CIR level rates are shared across the member ports of the LAG when configured in **distribute** mode.

This command does not have any affect when the port on which this **port-scheduler-policy** is configured is part of a LAG in **link** mode. Similarly when rates are configured as percent-active rates, the value of this object is irrelevant.

group

Syntax **group** *name* [**create**]
no group *name*

Context config>qos>port-scheduler-policy

Description This command defines a weighted scheduler group within a port scheduler policy.

The port scheduler policy defines a set of eight priority levels. The weighted scheduler group allows for the application of a scheduling weight to groups of child queues competing at the same priority level of the port scheduler policy applied to a vport defined in the context of the egress of an Ethernet port or applied to the egress of an Ethernet port.

Up to eight groups can be defined within each port scheduler policy. One or more levels can map to the same group. A group has a rate and optionally a cir-rate and inherits the highest scheduling priority of its member levels. In essence, a group receives bandwidth from the port or from the vport and distributes it within the member levels of the group according to the weight of each level within the group.

Each priority level will compete for bandwidth within the group based on its weight under a congestion situation. If there is no congestion, a priority level can achieve up to its rate (cir-rate) worth of bandwidth.

Note that CLI will enforce that mapping of levels to a group are contiguous. In other words, a user would not be able to add priority level to group unless the resulting set of priority levels is contiguous.

The **no** form of the command removes the group from the port scheduler policy.

Parameters *name* — Specifies the name of the weighted scheduler group and can be up to 32 ASCII characters in length.

create — This keyword is mandatory when creating the specified group.

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*]
no percent-rate

Context

Context config>qos>port-scheduler-policy>group

Description The percent-rate command within the port scheduler policy group enables supports for a policer's PIR and CIR rate to be configured as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

If the parent arbiter rate changes after the policer is created, the policer's PIR and CIR rates will be recalculated based on the defined percentage value.

The **rate** and **percent-rate** commands override one another. If the current rate for a policer is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A policer's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

The **no** form of this command returns the queue to its default shaping rate and cir rate.

Parameters *pir-percent* — Specifies the policer's PIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values Percentage ranging from 0.01 to 100.00.

Default 100.00

cir cir-percent — The **cir** keyword is optional and when defined the required cir-percent CIR parameter expresses the policer's CIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values Percentage ranging from 0.00 to 100.00.

Default 100.00

rate

Syntax **rate** *pir-percent* [**cir** *cir-percent*]
no rate

Context config>qos>port-scheduler-policy>group

Description This command specifies the total bandwidth and the within-cir bandwidth allocated to a weighted scheduler group.

Parameters *pir-percent* — Specifies PIR rates.

Values kilobits-per-second: 1 — 100000000, max, Kbps

cir *cir-percent* **—** Specifies CIR rates.

Values 0 — 100000000, max, Kbps

level

Syntax **level** *priority-level* **rate** *pir-rate* [**cir** *cir-rat*e] **group** *name* [**weight** *weight*]
level *priority-level* **rate** *pir-rate* [**cir** *cir-rat*e]
no level *priority-level*

Context config>qos>port-scheduler-policy

Description This command configures an explicit within-cir bandwidth limit and a total bandwidth limit for each port scheduler's priority level. To understand how to set the level rate and CIR parameters, a basic understanding of the port level scheduler bandwidth allocation mechanism is required. The port scheduler takes all available bandwidth for the port or channel (after the max-rate and any port egress-rate limits have been accounted for) and offers it to each of the eight priority levels twice.

The first pass is called the within-cir pass and consists of providing the available port bandwidth to each of the 8 priority levels starting with level 8 and moving down to level 1. Each level takes the offered load and distributes it to all child members that have a port-parent cir-level equal to the current priority level. (Any

child with a cir-weight equal to 0 is skipped in this pass.) Each child may consume bandwidth up to the child's frame based within-cir offered load. The remaining available port bandwidth is then offered to the next lower priority level until level 1 is reached.

The second pass is called the above-cir pass and consists of providing the remaining available port bandwidth to each of the eight priority levels a second time. Again, each level takes the offered load and distributes it to all child members that have a port-parent level equal to the current priority level. Each child may consume bandwidth up to the remainder of the child's frame based offered load (some of the offered load may have been serviced during the within-cir pass). The remaining available port bandwidth is then offered to the next priority level until level 1 is again reached.

If the port scheduling policy is using the default orphan behavior (orphan-override has not been configured on the policy), the system then takes any remaining port bandwidth and allocates it to the orphan queues and scheduler on priority level 1. In a non-override orphan state, all orphans are attached to priority level 1 using a weight of 0. The 0 weight value causes the system to allocate bandwidth equally to all orphans based on each orphan queue or scheduler's ability to use the bandwidth. If the policy has an orphan-override configured, the orphans are handled based on the override commands parameters in a similar fashion to properly parented queues and schedulers.

The port scheduler priority level command rate keyword is used to optionally limit the total amount of bandwidth that is allocated to a priority level (total for the within-cir and above-cir passes). The cir keyword optionally limits the first pass bandwidth allocated to the priority level during the within-cir pass.

When executing the level command, at least one of the optional keywords, **rate** or **cir,** must be specified. If neither keyword is included, the command will fail.

If a previous explicit value for rate or cir exists when the level command is executed, and either rate or cir is omitted, the previous value for the parameter is overwritten by the default value and the previous value is lost.

The configured priority level rate limits may be overridden at the egress port or channel using the egress-scheduler-override level priority-level command. When a scheduler instance has an override defined for a priority level, both the rate and cir values are overridden even when one of them is not explicitly expressed in the override command. For instance, if the cir kilobits-per-second portion of the override is not expressed, the scheduler instance defaults to not having a CIR rate limit for the priority level even when the port scheduler policy has an explicit CIR limit defined.

Default **no level priority-level**

Parameters *priority-level* — Specifies to which priority level the level command pertains. Each of the eight levels is represented by an integer value of 1 to 8, with 8 being the highest priority level.

Values 1 — 8 (8 is the highest priority)

rate *pir-rate* — Specifies the total bandwidth limits allocated to priority-level.

Values 1 — 40000000 (Kilobits per second (1,000 bits per second))

cir *cir-rate* — The cir-rate specified limits the total bandwidth allocated in the within-cir distribution pass to priority-level. When cir is not specified, all the available port or channel bandwidth may be allocated to the specified priority level during the within-cir pass.

Values 1 — 40000000 (Kilobits per second (1,000 bits per second))

The value given for kilobits-per-second is expressed in kilobits-per-second on a base 10 scale that is usual for line rate calculations. If a value of 1 is given, the result is 1000 bits per second (as opposed to a base 2 interpretation that would be 1024 bits per second).

group *name* **—** specifies the existing group which specifies the weighted scheduler group this level maps to.

weight *weight* **—** Specifies and integer which specifies the weight of the level within this weighted scheduler group.

Values 1 — 100

Default 1

max-rate

Syntax **max-rate** *pir-rate*
max-rate percent *percent-rate*
no max-rate

Context config>qos>port-scheduler-policy

Description This command defines an explicit maximum frame based bandwidth limit for the port scheduler policies scheduler context. By default, once a scheduler policy is associated with a port or channel, the instance of the scheduler on the port automatically limit the bandwidth to the lesser of port or channel line rate and a possible egress-rate value (for Ethernet ports). If a max-rate is defined that is smaller than the port or channel rate, the expressed kilobits-per-second value is used instead. The max-rate command is another way to sub-rate the port or channel.

The max-rate command may be executed at anytime for an existing port-scheduler-policy. When a new max-rate is given for a policy, the system evaluates all instances of the policy to see if the configured rate is smaller than the available port or channel bandwidth. If the rate is smaller and the maximum rate is not currently overridden on the scheduler instance, the scheduler instance is updated with the new maximum rate value.

The max-rate value defined in the policy may be overridden on each scheduler instance. If the maximum rate is explicitly defined as an override on a port or channel, the policies max-rate value has no effect.

The **no** form of this command removes an explicit rate value from the port scheduler policy. Once removed, all instances of the scheduler policy on egress ports or channel are allowed to run at the available line rate unless the instance has a max-rate override in place.

Parameters *pir-rate* — Specifies the PIR rate.

Values 1 — 800000000, max, in Kbps

percent *percent-rate* **—** Specifies the percent rate.

Values 0.01 — 100.00

orphan-override

Syntax **orphan-override** [**level** *priority-level*] [**weight** *percent*] [**cir-level** *priority-level*] [**cir-weight** *cir-weight*]
no orphan-override

Context config>qos>port-scheduler-policy

Description This command override the default orphan behavior for port schedulers created using the port scheduler policy. The default orphan behavior is to give all orphan queues and schedulers bandwidth after all other properly parented queues and schedulers. Orphans by default do not receive any within-cir bandwidth and receive above-cir bandwidth after priority levels 8 through 1 have been allocated. The orphan-override command accepts the same parameters as the port-parent command in the SAP egress and network queue policy contexts. The defined parameters are used as a default port-parent association for any queue or scheduler on the port that the port scheduler policy is applied.

Orphan queues and schedulers are identified as:

- Any queue or scheduler that does not have a port-parent or parent command applied
- Any queue that has a parent command applied, but the specified scheduler name does not exist on the queue's SAP, MSS or SLA Profile instance.

A queue or scheduler may be properly parented to an upper level scheduler, but that scheduler may be orphaned. In this case, the queue or scheduler receives bandwidth from its parent scheduler based on the parent schedulers ability to receive bandwidth as an orphan.

Within-CIR Priority Level Parameters

The within-cir parameters define which port priority level the orphan queues and schedulers should be associated with when receiving bandwidth for the queue or schedulers within-cir offered load. The within-cir offered load is the amount of bandwidth the queue or schedulers could use that is equal to or less than its defined or summed CIR value. The summed value is only valid on schedulers and is the sum of the within-cir offered loads of the children attached to the scheduler. The parameters that control within-cir bandwidth allocation for orphans are the orphan-override commands cir-level and cir-weight keywords. The cir-level keyword defines the port priority level that the scheduler or queue uses to receive bandwidth for its within-cir offered load. The cir-weight is used when multiple queues or schedulers exist at the same port priority level for within-cir bandwidth. The weight value defines the relative ratio that is used to distribute bandwidth at the priority level when more within-cir offered load exists than the port priority level has bandwidth.

A cir-weight equal to zero (the default value) has special meaning and informs the system that the orphan queues and schedulers do not receive bandwidth from the within-cir distribution. Instead all bandwidth for the orphan queues and schedulers must be allocated from the port scheduler's above-cir pass.

Above-CIR Priority Level Parameters

The above-cir parameters define which port priority level the orphan queues and schedulers should be associated with when receiving bandwidth for the queue or schedulers above-cir offered load. The above-cir offered load is the amount of bandwidth the queue or schedulers could use that is equal to or less than its defined PIR value (based on the queue or schedulers rate command) less any bandwidth that was given to

the queue or scheduler during the above-cir scheduler pass. The parameters that control above-cir bandwidth allocation for orphans are the orphan-override commands level and weight keywords. The level keyword defines the port priority level that the scheduler or queue uses to receive bandwidth for its above-cir offered load. The weight is used when multiple queues or schedulers exist at the same port priority level for above-cir bandwidth. The weight value defines the relative ratio that is used to distribute bandwidth at the priority level when more above-cir offered load exists than the port priority level has bandwidth.

The **no** form of the command removes the orphan override port parent association for the orphan queues and schedulers on port schedulers created with the port scheduler policy. Any orphan queues and schedulers on a port associated with the port scheduler policy will revert to default orphan behavior.

Parameters **level** *priority-level* — Defines the port priority the orphan queues and schedulers will use to receive bandwidth for its above-cir offered-load.

Values 1 — 8 (8 is the highest priority)

Default 1

weight *percent* **—** Defines the weight the orphan queues and schedulers will use in the above-cir port priority level (defined by the level parameter).

Values 1 — 100

Default 1

cir-level *priority-level* **—** Defines the port priority the orphan queues and schedulers will use to receive bandwidth for its within-cir offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the orphan queues and schedulers do not receive bandwidth during the port scheduler's within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 1 — 8 (8 is the highest level)

cir-weight *cir-weight* **—** Defines the weight the orphan queues and schedulers will use in the within-cir port priority level (defined by the cir-level parameter). When the cir-weight parameter is set to a value of 0 (the default value), the orphan queues and schedulers do not receive bandwidth during the port scheduler's within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 1 — 100 (100 is the highest weight)

Show Commands

scheduler-policy

Syntax **scheduler-policy** *scheduler-name* [**association | sap-ingress** *policy-id* **| sap-egress** *policy-id*]

Context show>qos

Description Use this command to display scheduler policy information.

Parameters *scheduler-name* — The name of a scheduler configured in the **config>qos>scheduler-policy** context.

association — Display the associations related to the specified scheduler name.

sap-ingress *policy-id* **—** Specify the SAP ingress QoS policy information.

sap-egress *policy-id* **—** Specify the SAP egress QoS policy information.

Output **Customer Scheduler-Policy Output —** The following table describes the customer scheduler hierarchy fields.

Table 39: Show QoS Scheduler-Policy Output Fields

Label	Description
Policy-Name	Specifies the scheduler policy name.
Description	A text string that helps identify the policy's context in the configuration file.
Tier	Specifies the level of hierarchy that a group of schedulers are associated with.
Scheduler	Specifies the scheduler name.
Lvl/Wt	Specifies the priority level of the scheduler when compared to other child schedulers and queues vying for bandwidth on the parent schedulers during the 'above CIR' distribution phase of bandwidth allocation. Weight defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same level.
Cir Lvl/Wt	Specifies the level of hierarchy when compared to other schedulers and queues when vying for bandwidth on the parent scheduler. Weight defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent scheduler.

Table 39: Show QoS Scheduler-Policy Output Fields (Continued)

Label	Description
PIR	Specifies the PIR rate.
CIR	Specifies the CIR rate.
Parent	Specifies the parent scheduler that governs the available bandwidth given the queue aside from the queue's PIR setting.
Service-Id	The ID that uniquely identifies the policy.
Customer-Id	The ID that uniquely identifies the customer.
SAP	Specifies the Service Access Point (SAP) within the service where the policy is applied.
Multi Service Site	Specifies the multi-service site name.
Orphan Queues	Specifies the number of queues in an orphaned state.
Hierarchy	Displays the scheduler policy tree structure.

Sample Output

```
A:ALA-12# show qos scheduler-policy SLA1
===============================================================================
QoS Scheduler Policy
===============================================================================
Policy-Name    : SLA1
Description    : NetworkControl(3), Voice(2) and NonVoice(1) have strict priorities
-------------------------------------------------------------------------------
Tier/Scheduler                     Lvl/Wt     PIR       Parent
                                   CIR Lvl/Wt CIR
-------------------------------------------------------------------------------
1 All_traffic                      1/1        11000     None
                                   -/-        max
2 NetworkControl                   3/1        100       All_traffic
                                   3/-        max
2 NonVoice                         1/1        11000     All_traffic
                                   1/-        max
2 Voice                            2/1        5500      All_traffic
                                   2/-        max
3 Internet_be                      1/1        max       NonVoice
                                   1/-        max
3 Internet_priority                2/1        max       NonVoice
                                   2/-        max
3 Internet_voice                   1/1        max       Voice
                                   -/-        max
3 VPN_be                           1/1        max       NonVoice
                                   1/-        max
3 VPN_nc                           1/1        100       NetworkControl
                                   -/-        36
3 VPN_priority                     2/1        max       NonVoice
```

```
                                    2/-         max
3 VPN_reserved                      3/1         max       NonVoice
                                    3/-         max
3 VPN_video                         5/1         1500      NonVoice
                                    5/-         1500
3 VPN_voice                         1/1         2500      Voice
                                    -/-         2500
===============================================================================
A:ALA-12#
A:ALA-12# show qos scheduler-policy SLA1 association
===============================================================================
QoS Scheduler Policy
===============================================================================
Policy-Name    : SLA1
Description    : NetworkControl(3), Voice(2) and NonVoice(1) have strict priorities
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Service-Id     : 6000 (Epipe)                Customer-Id  : 274
 - SAP : 1/1/3.1:0 (Egress)
Service-Id     : 7000 (VPLS)                 Customer-Id  : 7
 - SAP : 1/1/5:0 (Egress)
 - Multi Service Site : west  (Ingress)
===============================================================================
A:ALA-12#


A:ALA-12# show qos scheduler-policy SLA1 sap-ingress 100
===============================================================================
Compatibility : Scheduler Policy SLA1 & Sap Ingress 100
===============================================================================
Orphan Queues :
None Found

Hierarchy     :

Root
|
|---(S) : All_traffic
|   |
|   |---(S) : NetworkControl
|   |   |
|   |   |---(S) : VPN_nc
|   |   |   |
|   |   |   |---(Q) : 17
|   |   |   |
|   |   |   |---(Q) : 27
|   |
|   |---(S) : NonVoice
|   |   |
|   |   |---(S) : Internet_be
|   |   |
|   |   |---(S) : Internet_priority
|   |   |
|   |   |---(S) : VPN_be
|   |   |   |
|   |   |   |---(Q) : 10
|   |   |   |
|   |   |   |---(Q) : 20
```

```
|   |   |
|   |   |---(S) : VPN_priority
|   |   |   |
|   |   |   |---(Q) : 12
|   |   |   |
|   |   |   |---(Q) : 22
|   |   |
|   |   |---(S) : VPN_reserved
|   |   |   |
|   |   |   |---(Q) : 13
|   |   |   |
|   |   |   |---(Q) : 23
|   |   |
|   |   |---(S) : VPN_video
|   |   |   |
|   |   |   |---(Q) : 15
|   |   |   |
|   |   |   |---(Q) : 25
|   |
|   |---(S) : Voice
|   |   |
|   |   |---(S) : Internet_voice
|   |   |
|   |   |---(S) : VPN_voice
|   |   |   |
|   |   |   |---(Q) : 16
|   |   |   |
|   |   |   |---(Q) : 26
|
|---(Q) : 1
|
|---(Q) : 2
===============================================================================
A:ALA-12#

A:ALA-12# show qos scheduler-policy SLA1 sap-egress 101
===============================================================================
Compatibility : Scheduler Policy SLA1 & Sap Egress 101
===============================================================================
Orphan Queues :

None Found

Hierarchy     :

Root
|
|---(S) : All_traffic
|   |
|   |---(S) : NetworkControl
|   |   |
|   |   |---(S) : VPN_nc
|   |
|   |---(S) : NonVoice
|   |   |
|   |   |---(S) : Internet_be
|   |   |
|   |   |---(S) : Internet_priority
|   |   |
```

```
|   |   |---(S) : VPN_be
|   |   |
|   |   |---(S) : VPN_priority
|   |   |
|   |   |---(S) : VPN_reserved
|   |   |
|   |   |---(S) : VPN_video
|   |
|   |---(S) : Voice
|   |   |
|   |   |---(S) : Internet_voice
|   |   |
|   |   |---(S) : VPN_voice
===============================================================================
A:ALA-12#
```

scheduler-hierarchy customer

Syntax **scheduler-hierarchy customer** *customer-id* **site** *customer-site-name* [**scheduler** *scheduler-name*] [**ingress** | **egress**] [**detail**]

Context show>qos

Description This command displays the scheduler hierarchy per customer multi-service-site.

Parameters *customer customer-id* — Specifies the ID number associated with a particular customer.

Values 1 — 2147483647

site *customer-site-name* — The unique name customer site name.

scheduler *scheduler-name* **—** The unique scheduler name created in the context of the scheduler policy.

ingress — Displays ingress SAP customer scheduler stats.

egress — Displays egress SAP customer scheduler stats.

detail — Displays detailed information.

Output **Show QoS Scheduler-Hierarchy Customer Output —** The following table describes the customer scheduler hierarchy fields.

Label	Description
Legend	Admin CIR/PIR: Specifies the configured value of CIR/PIR. Assigned CIR/PIR: Specifies the PIR/CIR rate given to a member by that parent level. Offered CIR/PIR: Specifies the offered load on that member. Consumed CIR/PIR: Specifies the amount of scheduler bandwidth used by this member.

Label	Description (Continued)
Lvl/Wt	Specifies the priority level of the scheduler when compared to other child schedulers and queues vying for bandwidth on the parent schedulers during the 'above CIR' distribution phase of bandwidth allocation. Weight defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same level.
Cir Lvl/Wt	Specifies the level of hierarchy when compared to other schedulers and queues when vying for bandwidth on the parent scheduler. Weight defines the relative weight of this queue as compared to other child schedulers and queues while vying for bandwidth on the parent scheduler.
PIR	Specifies the PIR rate.
CIR	Specifies the CIR rate.
Parent	Specifies the parent scheduler that governs the available bandwidth given the queue aside from the queue's PIR setting.
Service-Id	The ID that uniquely identifies the policy.
Customer-Id	The ID that uniquely identifies the customer.
SAP	Specifies the Service Access Point (SAP) within the service where the policy is applied.
Multi Service Site	Specifies the multi-service site name.
Orphan Queues	Specifies the number of queues in an orphaned state.
Hierarchy	Displays the scheduler policy tree structure.

Sample Output

```
A:D# show qos scheduler-hierarchy customer 1 site bc
===============================================================================
Scheduler Hierarchy - Customer 1 MSS bc
===============================================================================
Root (Ing)
| slot(1)
|--(S) : gp
Root (Egr)
| slot(1)
|--(S) : gp
| |
| |--(S) : pb
| | |
| | |--(S) : pbs
```

```
| |
| |--(S) : mb
| | |
| | |--(S) : mbs
|
|--(S) : rb
| |
| |--(S) : rbs
===============================================================================
A:D#
```

scheduler-hierarchy port

Syntax **scheduler-hierarchy port** *port-id* [**detail**] **queue-group** *queue-group-name* [**scheduler** *scheduler-name*] [**ingress** | **egress**]
scheduler-hierarchy port *port-id* [**detail**]

Context show>qos

Description This command displays scheduler hierarchy information per port.

Parameters *port-id* — Specifies the port ID in the slot/mda/port[.channel] format.

detail — Displays detailed information.

queue-group *queue-group-name* **—** Displays information about the specified queue group on the port.

scheduler *scheduler-name* **—** Displays information about the specified scheduler policy on the port.

ingress — Specifies to display ingress queue group information.

egress — Specifies to display egress queue group information.

Output **Show QoS Scheduler-Hierarchy Port Output —** The following table describes port scheduler hierarchy fields.

Table 40: Show QoS Schedule-Hierarchy Port Output Fields

Label	Description
S	Displays the scheduler name.
Q	Displays the queue ID and information.
Admin CIR/PIR:	Specifies the configured value of CIR/PIR.
Assigned CIR/PIR:	Specifies the on-the-wire PIR/CIR rate given to a member by that parent level.
Offered CIR/PIR:	Specifies the on-the-wire offered load on that member.
Consumed CIR/PIR:	Specifies the amount of scheduler bandwidth used by this member.

Sample Output

```
*A:Dut-R# show qos scheduler-hierarchy port 1/2/1 detail
===============================================================================
Scheduler Hierarchy - Port 1/2/1
===============================================================================
Port-scheduler-policy p1
    Port Bandwidth : 10000000   Max Rate : max
    Consumed : 0          Offered : 0
[Within CIR Level 8]
    Rate : max
    Consumed : 0          Offered : 0

[Within CIR Level 7]
    Rate : max
    Consumed : 0          Offered : 0

[Within CIR Level 6]
    Rate : max
    Consumed : 0          Offered : 0

    (Q) : 2->1/2/1:1->3
    Assigned : 768        Offered : 0
    Consumed : 0
    Weight   : 0

[Within CIR Level 5]
    Rate : max
    Consumed : 0          Offered : 0

[Within CIR Level 4]
    Rate : max
    Consumed : 0          Offered : 0

[Within CIR Level 3]
    Rate : max
    Consumed : 0          Offered : 0

[Within CIR Level 2]
    Rate : max
    Consumed : 0          Offered : 0

    (S) voip(SAP 1/2/1:1)
    Assigned : 0          Offered : 0
    Consumed : 0
    Weight   : 40

    (S) all(SAP 1/2/1:1)
    Assigned : 19000      Offered : 0
    Consumed : 0
    Weight   : 50

[Within CIR Level 1]
    Rate : max
    Consumed : 0          Offered : 0

[Within CIR Level 0]
    Rate : 0
    Consumed : 0          Offered : 0
```

```
[Above CIR Level 8]
    Rate : max
    Consumed : 0          Offered : 0

[Above CIR Level 7]
    Rate : max
    Consumed : 0          Offered : 0

[Above CIR Level 6]
    Rate : max
    Consumed : 0          Offered : 0

[Above CIR Level 5]
    Rate : max
    Consumed : 0          Offered : 0

[Above CIR Level 4]
    Rate : max
    Consumed : 0          Offered : 0

[Above CIR Level 3]
    Rate : max
    Consumed : 0          Offered : 0

[Above CIR Level 2]
    Rate : max
    Consumed : 0          Offered : 0

    (S) voip(SAP 1/2/1:1)
    Assigned : 10000000   Offered : 0
    Consumed : 0
    Weight   : 30

    (S) all(SAP 1/2/1:1)
    Assigned : 960000     Offered : 0
    Consumed : 0
    Weight   : 50

[Above CIR Level 1]
    Rate : max
    Consumed : 0          Offered : 0

    (Q) : 2->1/2/1:1->3
    Assigned : 786        Offered : 0
    Consumed : 0
    Weight   : 1

===============================================================================
*A:Dut-R#
```

scheduler-hierarchy sap

Syntax **scheduler-hierarchy sap** *sap-id* [**scheduler** *scheduler-name*] [**ingress** | **egress**] [**detail**]

Context show>qos

Description This command displays the scheduler hierarchy per SAP.

Parameters **sap** *sap-id* **—** Specifies the SAP assigned to the service.

Values: *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id*][:*vpi/vci*|*vpi*| *vpi1.vpi2*]	
frame	[*port-id* | *aps-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
cem	*slot/mda/port.channel*	
ima-grp	[*bundle-id*[:vpi/vci|vpi|*vpi1.vpi2*]	
port-id	*slot/mda/port*[.*channel*]	
bundle-id	bundle-*type-slot/mda.bundle-num*	
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 336
bpgrp-id	bpgrp-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima, ppp
	bpgrp-num	1 — 256
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 16
lag-id	lag-*id*	
	lag	keyword
	id	1 — 64
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI: 0 — 4095	
	UNI: 0 — 255	
vci	1, 2, 5 — 65535	
dlci	16 — 1022	
ipsec-id	ipsec-*id*.[private | public]:*tag*	
	ipsec	keyword
	id	1 — 4
	tag	0 — 4094

scheduler *scheduler-name* **—** The unique scheduler name created in the context of the scheduler policy

ingress — The keyword to display ingress SAP scheduler stats.

egress — The keyword to display egress SAP scheduler stats.

detail — Displays detailed information.

Output **Show Qos Scheduler-Hierarchy SAP Output —** The following table describes the SAP scheduler hierarchy fields.

Table 41: Show QoS Scheduler-Hierarchy SAP Output Fields

Label	Description
Legend	Admin CIR/PIR: Specifies the configured value of CIR/PIR. Assigned CIR/PIR: Specifies the PIR/CIR rate given to a member by that parent level. Offered CIR/PIR: Specifies the offered load on that member. Consumed CIR/PIR: Specifies the amount of scheduler bandwidth used by this member.
PIR	Specifies the PIR rate.
CIR	Specifies the CIR rate.
S	Displays the scheduler name.
Q	Displays the queue ID and information.

Sample Output

```
*A:Dut-R# show qos scheduler-hierarchy sap 1/2/1:1 ingress detail
===============================================================================
Scheduler Hierarchy - Sap 1/2/1:1
===============================================================================
Legend :
(*) real-time dynamic value
(w) Wire rates
-------------------------------------------------------------------------------
Root (Ing)
| slot(1)
|--(S) : tplay
|   |    AdminPIR:960000     AdminCIR:960000(sum)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0          Offered:0
|   |    Consumed:0
|   |
|   |    [Above CIR Level 0 Weight 0]
|   |    Assigned:0          Offered:0
|   |    Consumed:0
|   |
|   |    TotalConsumed:0
|   |    OperPIR:960000
|   |
|   |    [As Parent]
|   |    Rate:960000
|   |    ConsumedByChildren:960000
|   |
|   |--(S) : voice
|   |   |    AdminPIR:max        AdminCIR:max(sum)
```

```
|   |   |
|   |   |    [Within CIR Level 6 Weight 1]
|   |   |    Assigned:960000     Offered:120000
|   |   |    Consumed:120000
|   |   |
|   |   |    [Above CIR Level 1 Weight 1]
|   |   |    Assigned:960000     Offered:120000
|   |   |    Consumed:0
|   |   |
|   |   |    TotalConsumed:120000
|   |   |    OperPIR:960000
|   |   |
|   |   |    [As Parent]
|   |   |    Rate:960000
|   |   |    ConsumedByChildren:120000
|   |   |
|   |   |--(S) : AccessIngress:2->1/2/1:1->3
|   |   |   |    AdminPIR:max        AdminCIR:max(sum)
|   |   |   |
|   |   |   |    [Within CIR Level 0 Weight 1]
|   |   |   |    Assigned:960000     Offered:0
|   |   |   |    Consumed:0
|   |   |   |
|   |   |   |    [Above CIR Level 1 Weight 1]
|   |   |   |    Assigned:960000     Offered:120000
|   |   |   |    Consumed:120000
|   |   |   |
|   |   |   |    TotalConsumed:120000
|   |   |   |    OperPIR:960000
|   |   |   |
|   |   |   |    [As Parent]
|   |   |   |    OperPIR:960000      OperCIR:960000
|   |   |   |    ConsumedByChildren:120000
|   |   |   |
|   |   |   |--(Q) : 2->1/2/1:1->3 5/1
|   |   |   |   |    AdminPIR:10000000   AdminCIR:10000000
|   |   |   |   |    CBS:6144            MBS:12288
|   |   |   |   |    Depth:0             HiPrio:2048
|   |   |   |   |
|   |   |   |   |    [CIR]
|   |   |   |   |    Assigned:960000     Offered:120000
|   |   |   |   |    Consumed:120000
|   |   |   |   |
|   |   |   |   |    [PIR]
|   |   |   |   |    Assigned:960000     Offered:120000
|   |   |   |   |    Consumed:0
|   |   |   |   |
|   |   |   |   |    OperPIR:960000      OperCIR:960000
|   |   |   |   |
|   |   |   |   |--(Q) : 2->1/2/1:1->3 1/2
|   |   |   |   |   |    AdminPIR:10000000   AdminCIR:10000000
|   |   |   |   |   |    CBS:6144            MBS:12288
|   |   |   |   |   |    Depth:0             HiPrio:2048
|   |   |   |   |   |
|   |   |   |   |   |    [CIR]
|   |   |   |   |   |    Assigned:840000     Offered:0
|   |   |   |   |   |    Consumed:0
|   |   |   |   |   |
|   |   |   |   |   |    [PIR]
```

```
|   |   |   |   |    Assigned:840000     Offered:0
|   |   |   |   |    Consumed:0
|   |   |   |   |
|   |   |   |   |    OperPIR:840000      OperCIR:840000
|   |   |   |   |
|   |   |
|   |   |--(S) : vod
|   |   |   |    AdminPIR:max        AdminCIR:max(sum)
|   |   |   |
|   |   |   |    [Within CIR Level 2 Weight 75]
|   |   |   |    Assigned:840000     Offered:2400000
|   |   |   |    Consumed:840000
|   |   |   |
|   |   |   |    [Above CIR Level 2 Weight 75]
|   |   |   |    Assigned:840000     Offered:2400000
|   |   |   |    Consumed:0
|   |   |   |
|   |   |   |    TotalConsumed:840000
|   |   |   |    OperPIR:840000
|   |   |   |
|   |   |   |    [As Parent]
|   |   |   |    Rate:840000
|   |   |   |    ConsumedByChildren:840000
|   |   |   |
|   |   |   |--(S) : AccessIngress:2->1/2/1:1->2
|   |   |   |   |    AdminPIR:max        AdminCIR:max(sum)
|   |   |   |   |
|   |   |   |   |    [Within CIR Level 0 Weight 1]
|   |   |   |   |    Assigned:840000     Offered:0
|   |   |   |   |    Consumed:0
|   |   |   |   |
|   |   |   |   |    [Above CIR Level 1 Weight 1]
|   |   |   |   |    Assigned:840000     Offered:2400000
|   |   |   |   |    Consumed:840000
|   |   |   |   |
|   |   |   |   |    TotalConsumed:840000
|   |   |   |   |    OperPIR:840000
|   |   |   |   |
|   |   |   |   |    [As Parent]
|   |   |   |   |    OperPIR:840000      OperCIR:840000
|   |   |   |   |    ConsumedByChildren:840000
|   |   |   |   |
|   |   |   |   |--(Q) : 2->1/2/1:1->2 5/1
|   |   |   |   |   |    AdminPIR:10000000   AdminCIR:10000000
|   |   |   |   |   |    CBS:6144            MBS:12288
|   |   |   |   |   |    Depth:10236         HiPrio:2048
|   |   |   |   |   |
|   |   |   |   |   |    [CIR]
|   |   |   |   |   |    Assigned:840000     Offered:2400000
|   |   |   |   |   |    Consumed:840000
|   |   |   |   |   |
|   |   |   |   |   |    [PIR]
|   |   |   |   |   |    Assigned:840000     Offered:2400000
|   |   |   |   |   |    Consumed:0
|   |   |   |   |   |
|   |   |   |   |   |    OperPIR:840000      OperCIR:840000
|   |   |   |   |   |
|   |   |   |   |--(Q) : 2->1/2/1:1->2 1/2
|   |   |   |   |   |    AdminPIR:10000000   AdminCIR:10000000
```

```
|   |   |   |   |    CBS:6144            MBS:12288
|   |   |   |   |    Depth:0             HiPrio:2048
|   |   |   |   |
|   |   |   |   |    [CIR]
|   |   |   |   |    Assigned:420000     Offered:0
|   |   |   |   |    Consumed:0
|   |   |   |   |
|   |   |   |   |    [PIR]
|   |   |   |   |    Assigned:420000     Offered:0
|   |   |   |   |    Consumed:0
|   |   |   |   |
|   |   |   |   |    OperPIR:420000      OperCIR:420000
|   |   |   |   |
|   |   |
|   |--(S) : hsi
|   |   |    AdminPIR:max        AdminCIR:0(sum)
|   |   |
|   |   |
|   |   |    [Within CIR Level 2 Weight 5]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [Above CIR Level 1 Weight 1]
|   |   |    Assigned:0          Offered:961000
|   |   |    Consumed:0
|   |   |
|   |   |    TotalConsumed:0
|   |   |    OperPIR:0
|   |   |
|   |   |    [As Parent]
|   |   |    Rate:0
|   |   |    ConsumedByChildren:0
|   |   |
|   |   |--(S) : AccessIngress:2->1/2/1:1->1
|   |   |   |    AdminPIR:max        AdminCIR:0(sum)
|   |   |   |
|   |   |   |    [Within CIR Level 0 Weight 1]
|   |   |   |    Assigned:0          Offered:0
|   |   |   |    Consumed:0
|   |   |   |
|   |   |   |    [Above CIR Level 1 Weight 1]
|   |   |   |    Assigned:0          Offered:961000
|   |   |   |    Consumed:0
|   |   |   |
|   |   |   |    TotalConsumed:0
|   |   |   |    OperPIR:0
|   |   |   |
|   |   |   |    [As Parent]
|   |   |   |    OperPIR:0           OperCIR:0
|   |   |   |    ConsumedByChildren:0
|   |   |   |
|   |   |   |--(Q) : 2->1/2/1:1->1 5/1
|   |   |   |   |    AdminPIR:10000000   AdminCIR:0
|   |   |   |   |    CBS:0               MBS:0
|   |   |   |   |    Depth:0             HiPrio:0
|   |   |   |   |
|   |   |   |   |    [CIR]
|   |   |   |   |    Assigned:0          Offered:0
|   |   |   |   |    Consumed:0
```

```
|   |   |   |   |
|   |   |   |   |    [PIR]
|   |   |   |   |    Assigned:0          Offered:961000
|   |   |   |   |    Consumed:0
|   |   |   |   |
|   |   |   |   |    OperPIR:0           OperCIR:0
|   |   |   |
|   |   |   |--(Q) : 2->1/2/1:1->1 1/2
|   |   |   |   |    AdminPIR:10000000   AdminCIR:0
|   |   |   |   |    CBS:0               MBS:0
|   |   |   |   |    Depth:0             HiPrio:0
|   |   |   |   |
|   |   |   |   |    [CIR]
|   |   |   |   |    Assigned:0          Offered:0
|   |   |   |   |    Consumed:0
|   |   |   |   |
|   |   |   |   |    [PIR]
|   |   |   |   |    Assigned:0          Offered:0
|   |   |   |   |    Consumed:0
|   |   |   |   |
|   |   |   |   |    OperPIR:0           OperCIR:0
===============================================================================
*A:Dut-R#


*A:Dut-R# show qos scheduler-hierarchy sap 5/1/1:1 egress detail
===============================================================================
Scheduler Hierarchy - Sap 5/1/1:1
===============================================================================
Legend :
(*) real-time dynamic value
(w) Wire rates
-------------------------------------------------------------------------------
Root (Egr)
| slot(5)
|--(S) : tplay
|   |    AdminPIR:960000     AdminCIR:19768(sum)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0          Offered:0
|   |    Consumed:0
|   |
|   |    [Above CIR Level 0 Weight 0]
|   |    Assigned:0          Offered:0
|   |    Consumed:0
|   |
|   |    TotalConsumed:0
|   |    OperPIR:960000
|   |
|   |    [As Parent]
|   |    Rate:960000
|   |    ConsumedByChildren:19661
|   |
|   |
|   |--(S) : hsi
|   |   |    AdminPIR:max        AdminCIR:3000(sum)
|   |   |
|   |   |    [Within CIR Level 2 Weight 5]
|   |   |    Assigned:3000       Offered:3000
```

```
|   |   |    Consumed:3000
|   |   |
|   |   |    [Above CIR Level 1 Weight 1]
|   |   |    Assigned:946339     Offered:6000
|   |   |    Consumed:3000
|   |   |
|   |   |    TotalConsumed:6000
|   |   |    OperPIR:946339
|   |   |
|   |   |    [As Parent]
|   |   |    Rate:946339
|   |   |    ConsumedByChildren:6000
|   |   |
|   |   |--(Q) : 2->5/1/1:1->1
|   |   |   |    AdminPIR:6000       AdminCIR:3000
|   |   |   |    CBS:4               MBS:64
|   |   |   |    Depth:56            HiPrio:8
|   |   |   |
|   |   |   |    [Within CIR Level 0 Weight 1]
|   |   |   |    Assigned:3000       Offered:0
|   |   |   |    Consumed:0
|   |   |   |
|   |   |   |    [Above CIR Level 1 Weight 1]
|   |   |   |    Assigned:6000       Offered:6000
|   |   |   |    Consumed:6000
|   |   |   |
|   |   |   |    TotalConsumed:6000
|   |   |   |    OperPIR:6000        OperCIR:3000
|   |   |
|   |--(S) : vod
|   |   |    AdminPIR:max        AdminCIR:16000(sum)
|   |   |
|   |   |    [Within CIR Level 2 Weight 75]
|   |   |    Assigned:16000      Offered:13100
|   |   |    Consumed:13100
|   |   |
|   |   |    [Above CIR Level 2 Weight 75]
|   |   |    Assigned:956439     Offered:13100
|   |   |    Consumed:0
|   |   |
|   |   |    TotalConsumed:13100
|   |   |    OperPIR:956439
|   |   |
|   |   |    [As Parent]
|   |   |    Rate:956439
|   |   |    ConsumedByChildren:13100
|   |   |
|   |   |--(Q) : 2->5/1/1:1->2
|   |   |   |    AdminPIR:20000      AdminCIR:16000
|   |   |   |    CBS:20              MBS:64
|   |   |   |    Depth:0             HiPrio:8
|   |   |   |
|   |   |   |    [Within CIR Level 0 Weight 1]
|   |   |   |    Assigned:16000      Offered:0
|   |   |   |    Consumed:0
|   |   |   |
|   |   |   |    [Above CIR Level 1 Weight 1]
|   |   |   |    Assigned:20000      Offered:13100
|   |   |   |    Consumed:13100
```

```
|   |   |   |
|   |   |   |    TotalConsumed:13100
|   |   |   |    OperPIR:20000       OperCIR:16000
|   |   |
|   |--(S) : voice
|   |   |    AdminPIR:max        AdminCIR:768(sum)
|   |   |
|   |   |    [Within CIR Level 6 Weight 1]
|   |   |    Assigned:768        Offered:561
|   |   |    Consumed:561
|   |   |
|   |   |    [Above CIR Level 1 Weight 1]
|   |   |    Assigned:940900     Offered:561
|   |   |    Consumed:0
|   |   |
|   |   |    TotalConsumed:561
|   |   |    OperPIR:940900
|   |   |
|   |   |    [As Parent]
|   |   |    Rate:940900
|   |   |    ConsumedByChildren:561
|   |   |
|   |   |--(Q) : 2->5/1/1:1->3
|   |   |   |    AdminPIR:786        AdminCIR:768
|   |   |   |    CBS:8               MBS:64
|   |   |   |    Depth:0             HiPrio:8
|   |   |   |
|   |   |   |    [Within CIR Level 0 Weight 1]
|   |   |   |    Assigned:768        Offered:0
|   |   |   |    Consumed:0
|   |   |   |
|   |   |   |    [Above CIR Level 1 Weight 1]
|   |   |   |    Assigned:786        Offered:561
|   |   |   |    Consumed:561
|   |   |   |
|   |   |   |    TotalConsumed:561
|   |   |   |    OperPIR:784         OperCIR:768
===============================================================================
*A:Dut-R#
```

scheduler-hierarchy subscriber

Syntax **scheduler-hierarchy subscriber** *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress** | **egress**] [**detail**]

Context show>qos

Description This command displays the scheduler hierarchy per subscriber.

Parameters **subscriber *sub-ident-string* —** Displays the subscriber identification policy name.

scheduler *scheduler-name* — Displays the scheduler name.

ingress — Displays ingress SAP subscriber scheduler stats.

egress — Displays egress SAP subscriber scheduler stats.

detail — Displays detailed information.

Output **Show QoS Scheduler-Hierarchy Subscriber Output —** The following table describes the QoS scheduler hierarchy subscriber fields.

Table 42: Show QoS Scheduler-Hierarchy Subscriber Output Fields

Label	Description
Legend	Admin CIR/PIR: Specifies the configured value of CIR/PIR. Assigned CIR/PIR: Specifies the PIR/CIR rate given to a member by that parent level. Offered CIR/PIR: Specifies the offered load on that member. Consumed CIR/PIR: Specifies the amount of scheduler bandwidth used by this member.
PIR	Specifies the PIR rate.
CIR	Specifies the CIR rate.
S	Displays the scheduler name.
Q	Displays the queue ID and information.

Sample Output

```
A:D# show qos scheduler-hierarchy subscriber RoutedCoHost1
===============================================================================
Scheduler Hierarchy - Subscriber RoutedCoHost1
===============================================================================
Root (Ing)
| slot(1)
|--(S) : grandpa
|   |
|   |--(S) : AccessIngress:Sub=1:1 200->1/2/5:1->8
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->8 1/2
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->8 2/1
|   |   |
|   |
|   |--(S) : AccessIngress:Sub=1:1 200->1/2/5:1->7
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->7 1/2
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->7 2/1
|   |   |
|   |
|   |--(S) : AccessIngress:Sub=1:1 200->1/2/5:1->6
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->6 1/2
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->6 2/1
|   |   |
```

```
|   |
|   |--(S) : AccessIngress:Sub=1:1 200->1/2/5:1->5
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->5 1/2
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->5 2/1
|   |   |
|   |
|   |--(S) : AccessIngress:Sub=1:1 200->1/2/5:1->4
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->4 1/2
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->4 2/1
|   |   |
|   |
|   |--(S) : AccessIngress:Sub=1:1 200->1/2/5:1->3
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->3 1/2
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->3 2/1
|   |   |
|   |
|   |--(S) : AccessIngress:Sub=1:1 200->1/2/5:1->2
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->2 1/2
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->2 2/1
|   |   |
|   |
|   |--(S) : AccessIngress:Sub=1:1 200->1/2/5:1->1
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->1 1/2
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->1 2/1
|   |   |
```

```
Root (Egr)
| slot(1)
|--(S) : gp
|   |
|   |--(S) : pb
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->1
|   |   |
|   |   |--(S) : pbs
|   |   |   |
|   |   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->2
|   |   |   |
|   |
|   |--(S) : mb
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->3
|   |   |
|   |   |--(S) : mbs
|   |   |   |
|   |   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->4
|   |   |   |
|
```

```
|--(S) : rb
|   |
|   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->5
|   |
|   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->7
|   |
|   |--(S) : rbs
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->6
|   |   |
|   |   |--(Q) : Sub=RoutedCoHost1:adsl-hsi 200->1/2/5:1->8
|   |   |
```

```
show qos scheduler-hierarchy subscriber x detail
...

|--(Q) : Sub=hpolSub43:hpolSlaProf1 2000->2/1/5:2000->2  (Port 2/1/5)
|   |    AdminPIR:100000     AdminCIR:0
|   |    AvgFrmOv:0.00
|   |    AdminPIR:100000(w)  AdminCIR:0(w)
|   |    CBS:0          B    MBS:125952     B
|   |    Depth:0          B  HiPrio:15360      B
|   |    MaxAggRate:22032821(w) CurAggRate:0(w)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0(w)       Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |    [Above CIR Level 1 Weight 60]
|   |    Assigned:1000(w)    Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |    TotalConsumed:0
|   |    OperPIR:1000        OperCIR:0
...
```

scheduler-name

Syntax **scheduler-name** *scheduler-name*

Context show>qos

Description This command displays the scheduler policies using the specified scheduler.

Parameters *scheduler-name* — The name of a scheduler configured in the **config>qos>scheduler-policy>tier** context.

Sample Output

```
A:ALA-12# show qos scheduler-name NetworkControl
===============================================================
Scheduler : NetworkControl
===============================================================
Scheduler Policy   : SLA1
```

```
Scheduler Policy   : alpha
Scheduler Policy   : beta
===============================================================
A:ALA-12#
```

scheduler-stats customer

Syntax **scheduler-stats customer** *customer-id* **site** *customer-site-name* [**scheduler** *scheduler-name*] [**ingress | egress**]

Context show>qos

Description This command displays scheduler statistics information.

Parameters **customer** *customer-id* **—** Specifies the ID number associated with a particular customer.

Values 1 — 2147483647

site *customer-site-name* **—** The unique customer site name.

scheduler *scheduler-name* **—** The unique scheduler name created in the context of the scheduler policy

ingress — The keyword to display ingress SAP customer scheduler stats.

egress — The keyword to display egress SAP customer scheduler stats.

Output **Show QoS Scheduler-Stats Customer Output —** The following table describes the SAP scheduler-stats customer fields.

Table 43: Show QoS Scheduler-Stats Customer Output Fields

Label	Description
Scheduler	Displays the scheduler policy name.
Forwarded Packets	Displays the number of packets forwarded.
Forwarded Octets	Displays the number of octets forwarded.

Sample Output

```
A:ALA-12# show qos scheduler-stats customer 274 site west scheduler NetworkControl ingress
===============================================================================
Scheduler Stats
===============================================================================
Scheduler                         Forwarded Packets        Forwarded Octets
-------------------------------------------------------------------------------
NetworkControl                    0                        0
===============================================================================
A:ALA-12#
```

scheduler-stats sap

Syntax **scheduler-stats sap** *sap-id* [**scheduler** *scheduler-name*] [**ingress** | **egress**]

Context show>qos

Description Display the scheduler stats per SAP.

Parameters **sap** *sap-id* **—** The port number and encapsulation value used to identify the SAP.

Values: *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id*][:*vpi/vci*|*vpi*| *vpi1.vpi2*]	
frame	[*port-id* | *aps-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
cem	*slot/mda/port.channel*	
ima-grp	[*bundle-id*[:vpi/vci|vpi|*vpi1.vpi2*]	
port-id	*slot/mda/port*[.*channel*]	
bundle-id	bundle-*type-slot/mda.bundle-num*	
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 336
bpgrp-id	bpgrp-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima, ppp
	bpgrp-num	1 — 256
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 16
eth-tunnel	eth-tunnel-*id*[:*eth-tun-sap*-id]	
	id:	1 — 1024
	eth-tun-sap-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 64
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI: 0 — 4095	
	UNI: 0 — 255	
vci	1, 2, 5 — 65535	
dlci	16 — 1022	
ipsec-id	ipsec-*id*.[private | public]:*tag*	
	ipsec	keyword
	id	1 — 4
	tag	0 — 4094

scheduler *scheduler-name* **—** The name of an existing scheduler policy.

ingress — Display only the policy displayed on the ingress SAP.

egress — Display only the policy displayed on the egress SAP.

Output **Show QoS Scheduler-Stats SAP Output —** The following table describes the scheduler-stats SAP fields.

Table 44: Show QoS Scheduler-Stats SAP Output Fields

Label	Description
Scheduler	Displays the scheduler policy name.
Forwarded Packets	Displays the number of packets forwarded.
Forwarded Octet	Displays the number of octets forwarded.
Ingress Schedulers	Displays the egress scheduler name(s).
Egress Schedulers	Displays the ingress scheduler name(s).

Sample Output

```
A:ALA-12# show qos scheduler-stats sap 1/1/4.1:0
===============================================================================
Scheduler Stats
===============================================================================
Scheduler                          Forwarded Packets        Forwarded Octets
-------------------------------------------------------------------------------
Ingress Schedulers
All_traffic                        0                        0
NetworkControl                     0                        0
Egress Schedulers
All_traffic                        0                        0
Internet_be                        0                        0
Internet_priority                  0                        0
Internet_voice                     0                        0
NetworkControl                     0                        0
NonVoice                           0                        0
VPN_be                             0                        0
VPN_nc                             0                        0
VPN_priority                       0                        0
VPN_reserved                       0                        0
VPN_video                          0                        0
VPN_voice                          0                        0
Voice                              0                        0
===============================================================================
A:ALA-12#


A:ALA-12# show qos scheduler-stats sap 1/1/5:0 scheduler 1
===============================================================================
Scheduler Stats
===============================================================================
Scheduler                          Forwarded Packets        Forwarded Octets
-------------------------------------------------------------------------------
Ingress Schedulers
No Matching Entries.
Egress Schedulers
No Matching Entries.
```

```
===============================================================================
A:ALA-12#


A:ALA-12# show qos scheduler-stats sap 1/1/4.1:0 scheduler All_traffic
===============================================================================
Scheduler Stats
===============================================================================
Scheduler                          Forwarded Packets        Forwarded Octets
-------------------------------------------------------------------------------Ingress
Schedulers
All_traffic                        0                        0
Egress Schedulers
All_traffic                        0                        0
===============================================================================
A:ALA-12#
```

scheduler-stats subscriber

Syntax **scheduler-stats subscriber** *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress | egress**]

Context show>qos

Description This command displays scheduler statistics information.

Parameters **subscriber *sub-ident-string* —** Specifies an existing SLA profile string.

scheduler *scheduler-name* — Specifies an existing scheduler name.

ingress — Display only the policy displayed on ingress.

egress — Display only the policy displayed on egress.

Output **Show QoS Scheduler-Stats Subscriber Output —** The following table describes the QoS scheduler-stats subscriber fields.

Table 45: Show QoS Scheduler-Stats Subscriber Output Fields

Label	Description
Scheduler	Displays the scheduler policy name.
Forwarded Packets	Displays the number of packets forwarded.
Forwarded Octet	Displays the number of octets forwarded.

Sample Output

```
A:D# show qos scheduler-stats subscriber RoutedCoHost1
===============================================================================
Scheduler Stats
===============================================================================
Scheduler                          Forwarded Packets      Forwarded Octets
-------------------------------------------------------------------------------
```

```
Ingress Schedulers
gp                                  0                       0
Egress Schedulers
gp                                  0                       0
mb                                  0                       0
mbs                                 0                       0
pb                                  0                       0
pbs                                 0                       0
rb                                  0                       0
rbs                                 0                       0
===============================================================================
*A:D#
```

port-scheduler-policy

Syntax **port-scheduler-policy** [*port-scheduler-policy-name*] [**association**]
port-scheduler-policy *port-scheduler-policy-name* **network-policy** *network-queue-policy-name*
port-scheduler-policy *port-scheduler-policy-name* **sap-egress** *policy-id*
port-scheduler-policy *port-scheduler-policy-name* **scheduler-policy** *scheduler-policy-name*
port-scheduler-policy *port-scheduler-policy-name* **scheduler-policy** *scheduler-policy-name* **sap-egress** *policy-id*

Context show>qos

Description This command displays port-scheduler policy information

Parameters *port-scheduler-policy-name* — Displays information for the specified existing port scheduler policy.

association — Displays associations related to the specified port scheduler policy.

network-policy *network-queue-policy-name* **—** Displays information for the specified existing network queue policy.

sap-egress *policy-id* **—** Displays information for the specified existing SAP egress policy.

scheduler-policy *scheduler-policy-name* **—** Displays information for the specified existing scheduler policy.

Output **Show QoS Port Scheduler Output —** The following table describes the QoS port scheduler policy fields.

Label	Description
Policy Name	Displays the port scheduler policy name.
Max Rate	Displays the explicit maximum frame-based bandwidth limit of this port scheduler.
Lvlx PIR	Displays the total bandwidth limit, PIR, for the specified priority level.
Lvlx CIR	Displays the within-cir bandwidth limit for the specified priority level.
Orphan Lvl	Displays above-cir port priority of orphaned queues and scheduler.
Orphan Weight	Displays the weight of orphaned queues and schedulers that are above-cir.
Orphan CIR-Lvl	Displays the port priority of orphaned queues and schedulers that are within-cir.
Orphan CIR-Weight	Displays the weight of orphaned queues and schedulers that are within-cir.

Label	Description (Continued)
Associations	Displays associations related to the specified port scheduler policy.
Mode	Displays the port scheduler policy mode (STRICT, RR, WRR, WDRR).
Accounting	Displays whether the accounting mode is frame-based or packet-based
Last Changed	Displays the last time the configuration changed.
Queue #	Displays the weight of the queue if configured.

Sample Output

```
*A:Dut-R# show qos port-scheduler-policy p1
===============================================================================
QoS Port Scheduler Policy
===============================================================================
Policy-Name       : p1
Max Rate          : max                Last changed      : 05/21/2007 10:39:15
Lvl1 PIR          : max                Lvl1 CIR          : max
Lvl2 PIR          : max                Lvl2 CIR          : max
Lvl3 PIR          : max                Lvl3 CIR          : max
Lvl4 PIR          : max                Lvl4 CIR          : max
Lvl5 PIR          : max                Lvl5 CIR          : max
Lvl6 PIR          : max                Lvl6 CIR          : max
Lvl7 PIR          : max                Lvl7 CIR          : max
Lvl8 PIR          : max                Lvl8 CIR          : max
Orphan Lvl        : default            Orphan Weight     : default
Orphan CIR-Lvl    : default            Orphan CIR-Weight : default
===============================================================================
QoS Port Scheduler Policy
===============================================================================
Policy-Name       : p1
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
 - Port : 5/1/1
===============================================================================
*A:Dut-R#
```

Sample Output

The following configuration displays **dist-lag-rate-shared** and **percent-rate** for level, **group** and **max-rate** in a port-scheduler-policy

```
*B:vineshDut-A>config>qos>port-sched-plcy# info
----------------------------------------------
 dist-lag-rate-shared
  max-rate percent 30.00
```

```
 group "test" create
               percent-rate 20.00 cir 20.00
  exit
  level 1 percent-rate 10.00 percent-cir 10.00
  level 2 percent-rate 20.00 percent-cir 20.00
  level 3 percent-rate 30.00 percent-cir 30.00
  level 4 percent-rate 40.00 percent-cir 40.00
  level 5 percent-rate 50.00 percent-cir 50.00
  level 6 percent-rate 60.00 percent-cir 60.00
  level 7 percent-rate 70.00 percent-cir 70.00
  level 8 percent-rate 80.00 percent-cir 80.00
```

Overrides

```
*B:vineshDut-A>config>port# info
----------------------------------------------
        ethernet
            mode access
            egress-scheduler-policy "psp2"
            egress-scheduler-override create
                max-rate percent 50.00
                level 1 percent-rate 10.00 percent-cir 10.00
                level 2 percent-rate 20.00 percent-cir 20.00
                level 3 percent-rate 30.00 percent-cir 30.00
                level 4 percent-rate 40.00 percent-cir 40.00
                level 5 percent-rate 50.00 percent-cir 50.00
                level 6 percent-rate 60.00 percent-cir 60.00
                level 7 percent-rate 70.00 percent-cir 70.00
                level 8 percent-rate 80.00 percent-cir 80.00
            exit
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------
```

The following output shows a **port-scheduler-policy** showing Dist Lag Rate and percent parameters

```
*B:vineshDut-A>config>port# /show qos port-scheduler-policy "psp2"
===============================================================================
QoS Port Scheduler Policy
===============================================================================
Policy-Name       : psp2
Description       : (Not Specified)
Max Rate          : max                Max Rate Percent  : 30.00
Dist LAG Rate     : True               Last changed      : 07/16/2014 21:31:51
Group             : test
Group PIR         : max                Group CIR         : max
Group PIR Percent : 20.00              Group CIR Percent : 20.00

Lvl1 PIR          : max                Lvl1 CIR          : max
Lvl1 PIR Percent  : 10.00              Lvl1 CIR Percent  : 10.00
Lvl2 PIR          : max                Lvl2 CIR          : max
Lvl2 PIR Percent  : 20.00              Lvl2 CIR Percent  : 20.00
Lvl3 PIR          : max                Lvl3 CIR          : max
Lvl3 PIR Percent  : 30.00              Lvl3 CIR Percent  : 30.00
Lvl4 PIR          : max                Lvl4 CIR          : max
Lvl4 PIR Percent  : 40.00              Lvl4 CIR Percent  : 40.00
```

```
Lvl5 PIR          : max                Lvl5 CIR          : max
Lvl5 PIR Percent  : 50.00              Lvl5 CIR Percent  : 50.00
Lvl6 PIR          : max                Lvl6 CIR          : max
Lvl6 PIR Percent  : 60.00              Lvl6 CIR Percent  : 60.00
Lvl7 PIR          : max                Lvl7 CIR          : max
Lvl7 PIR Percent  : 70.00              Lvl7 CIR Percent  : 70.00
Lvl8 PIR          : max                Lvl8 CIR          : max
Lvl8 PIR Percent  : 80.00              Lvl8 CIR Percent  : 80.00
Orphan Lvl        : default            Orphan Weight     : default
Orphan CIR-Lvl    : default            Orphan CIR-Weight : default
===============================================================================
Part of show port  Output
-------------------------------------------------------------------------------
Egr Port Sched Override
-------------------------------------------------------------------------------
Max Rate           : max*                     Max Rate Percent : 50.00
Lvl1 PIR           : max*                     Lvl1 CIR         : max*
Lvl1 PIR Percent   : 10.00                    Lvl1 CIR Percent : 10.00
Lvl2 PIR           : max*                     Lvl2 CIR         : max*
Lvl2 PIR Percent   : 20.00                    Lvl2 CIR Percent : 20.00
Lvl3 PIR           : max*                     Lvl3 CIR         : max*
Lvl3 PIR Percent   : 30.00                    Lvl3 CIR Percent : 30.00
Lvl4 PIR           : max*                     Lvl4 CIR         : max*
Lvl4 PIR Percent   : 40.00                    Lvl4 CIR Percent : 40.00
Lvl5 PIR           : max*                     Lvl5 CIR         : max*
Lvl5 PIR Percent   : 50.00                    Lvl5 CIR Percent : 50.00
Lvl6 PIR           : max*                     Lvl6 CIR         : max*
Lvl6 PIR Percent   : 60.00                    Lvl6 CIR Percent : 60.00
Lvl7 PIR           : max*                     Lvl7 CIR         : max*
Lvl7 PIR Percent   : 70.00                    Lvl7 CIR Percent : 70.00
Lvl8 PIR           : max*                     Lvl8 CIR         : max*
Lvl8 PIR Percent   : 80.00                    Lvl8 CIR Percent : 80.00
* means the value is inherited
-------------------------------------------------------------------------------
```

Clear Commands

sap

Syntax **sap** *sap-id* [**scheduler** *scheduler-name*] [**ingress** | **egress**]

Context clear>qos>scheduler-stats

Description This command clears scheduler statistics.

Parameters *sap-id* — Specifies the SAP assigned to the service.

Values: *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id*][:*vpi/vci*|*vpi*| *vpi1.vpi2*]	
frame	[*port-id* | *aps-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
cem	*slot/mda/port.channel*	
ima-grp	[*bundle-id*[:vpi/vci|vpi|*vpi1.vpi2*]	
port-id	*slot/mda/port*[.*channel*]	
bundle-id	bundle-*type-slot/mda.bundle-num*	
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 336
bpgrp-id	bpgrp-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima, ppp
	bpgrp-num	1 — 256
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 16
lag-id	lag-*id*	
	lag	keyword
	id	1 — 64
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI: 0 — 4095	
	UNI: 0 — 255	
vci	1, 2, 5 — 65535	
dlci	16 — 1022	
ipsec-id	ipsec-*id*.[private | public]:*tag*	
	ipsec	keyword
	id	1 — 4
	tag	0 — 4094

scheduler-name — The name of the scheduler.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

ingress — Display only the policy displayed on the ingress SAP.

egress — Display only the policy displayed on the egress SAP.

Slope QoS Policies

In This Section

This section provides information to configure slope QoS policies using the command line interface.

Topics in this section include:

- Overview on page 606
- Basic Configurations on page 607
- Default Slope Policy Values on page 610
- Deleting QoS Policies on page 611

Overview

Default buffer pools exist (logically) at the port, MDA and node levels. Each physical port has three associated pool objects:

- Access ingress pool
- Access egress pool
- Network egress pool

Each MDA has three associated pool objects:

- Access egress pool
- Access ingress pool
- Network egress pool

The overall node has one associated pool object:

- Network ingress pool
- By default, each pool is associated with slope-policy default which disables the high-slope and low-slope parameters within the pool.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your router, refer to CLI Usage chapter in the Basic System Configuration Guide.

Basic Configurations

A basic slope QoS policy must conform to the following:

- Each slope policy must have a unique policy ID.
- High slope and low slope are shut down (default).
- Default values can be modified but parameters cannot be deleted.

Create a Slope QoS Policy

Configuring and applying slope policies is optional. If no slope policy is explicitly applied to a SAP or IP interface, a default slope policy is applied.

To create a new slope policy, define the following:

- A slope policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- The high slope for the high priority Random Early Detection (RED) slope graph.
- The low slope for the low priority Random Early Detection (RED) slope graph.
- The time average factor (TAF), a weighting exponent used to determine the portion of the shared buffer instantaneous utilization and shared buffer average utilization used to calculate the new shared buffer average utilization.

Use the following CLI syntax to configure a slope policy:

CLI Syntax:

```
config>qos#
 slope-policy name
    description description-string
    high-slope
       start-avg percent
       max-avg percent
       max-prob percent
       no shutdown
    low-slope
       start-avg percent
       max-avg percent
       max-prob percent
       no shutdown
    time-average-factor taf
```

The following displays the slope policy configuration:

```
ALA-7>config>qos# info
#------------------------------------------
echo "QoS Slope/Queue Policies Configuration"
#------------------------------------------
...
        slope-policy "slopePolicy1" create
            description "Test"
            high-slope
                no shutdown
            exit
            low-slope
                no shutdown
            exit
        exit
...
#------------------------------------------
ALA-7>config>qos#
```

Applying Slope Policies

Apply slope policies to the following entities:

- Global
- MDA
- MDA Ports

Global

Use the following CLI syntax to apply slope policies to network egress and ingress pools.

CLI Syntax: `config> card 1 mda 1 network ingress pool slope-policy name port`

MDA

The following CLI syntax examples may be used to apply slope policies to MDAs:

CLI Syntax:
```
config>card>mda>access>ingress>pool>slope-policy name
config>card>mda>network>egress>pool>slope-policy name
```

The following CLI syntax example configures the PPP multilink pool:

CLI Syntax: `config>card>mda>access>egress>pool>slope-policy name`

MDA Ports

The following CLI syntax examples may be used to apply slope policies to MDA ports:

CLI Syntax:
```
config>port>access>egress>pool>slope-policy name
config>port>network>egress>pool>slope-policy name
```

Default Slope Policy Values

The default access ingress and egress policies are identified as policy-id 1. The default policies cannot be edited or deleted. The following displays default policy parameters:

Table 46: Slope Policy Defaults

Field	Default
description	"Default slope policy"
high-slope	
shutdown	shutdown
start-age	70
max-avg	90
max-prob	80
low-slope	
shutdown	shutdown
start-age	50
max-avg	75
max-prob	80
time-average-factor	7

The following output displays the default configuration:

```
ALA-7>config>qos>slope-policy# info detail
----------------------------------------------
            description "Default slope policy."
            high-slope
                shutdown
                start-avg 70
                max-avg 90
                max-prob 80
            exit
            low-slope
                shutdown
                start-avg 50
                max-avg 75
                max-prob 80
            exit
            time-average-factor 7
----------------------------------------------
ALA-7>config>qos>slope-policy#
```

Deleting QoS Policies

A slope policy is associated by default with CMA and port access and network egress pools. A default policy may be replaced with a non-default policy, but a policy cannot be entirely removed from the configuration. When a non-default policy is removed, the policy association reverts to the default slope **policy** *policy-id* **default**. A QoS policy cannot be deleted until it is removed from all CMAs or ports where it is applied.

```
ALA-7>config>qos# no slope-policy slopePolicy1
MINOR: QOS #1902 Slope policy has references
ALA-7>config>qos#
```

Global

Use the following CLI syntax to remove slope policies from network egress and ingress pools.

CLI Syntax: `config> card 1 mda 1 network ingress pool` **`no`** `slope-policy name port`

CMA

The following CLI syntax examples can be used to remove slope policies from CMAs:

CLI Syntax:
```
config>card>mda>access>ingress>pool# no slope-policy name
config>card>mda>access>egress>pool# no slope-policy name
config>card>mda>network>egress>pool# no slope-policy name
```

MDA

The following CLI syntax examples can be used to remove slope policies from MDAs:

CLI Syntax:
```
config>card>mda>access>ingress>pool# no slope-policy name
config>card>mda>network>egress>pool# no slope-policy name
```

The following CLI syntax example configures the PPP multilink pool:

CLI Syntax: `config>card>mda>access>egress>pool#` **`no`** `slope-policy name`

MDA Ports

The following CLI syntax examples can be used to remove slope policies from MDA ports:

CLI Syntax: config>port>access>egress>pool# **no** slope-policy name
config>port>network>egress>pool# **no** slope-policy name

Remove a Policy from the QoS Configuration

To delete a slope policy, enter the following command:

CLI Syntax: config>qos# no slope-policy *policy-id*

Example: config>qos# no slope-policy slopePolicy1

Copying and Overwriting QoS Policies

You can copy an existing slope policy, rename it with a new policy ID value, or overwrite an existing policy ID. The `overwrite` option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos> copy {slope-policy}` *`source-policy-id dest-policy-id`* `[overwrite]`

The following output displays the copied policies:

```
ALA-7>config>qos# info
----------------------------------------------
...
        slope-policy "default" create
            description "Default slope policy."
            high-slope
                shutdown
                start-avg 70
                max-avg 90
                max-prob 80
            exit
            low-slope
                shutdown
                start-avg 50
                max-avg 75
                max-prob 80
            exit
            time-average-factor 7
        exit
        slope-policy "slopePolicy1" create
            description "Default slope policy."
            high-slope
                shutdown
                start-avg 70
                max-avg 90
                max-prob 80
            exit
            low-slope
                shutdown
                start-avg 50
                max-avg 75
                max-prob 80
            exit
            time-average-factor 7
        exit
        slope-policy "slopePolicy2" create
            description "Default slope policy."
            high-slope
                shutdown
                start-avg 70
                max-avg 90
                max-prob 80
            exit
            low-slope
```

```
                shutdown
                start-avg 50
                max-avg 75
                max-prob 80
            exit
            time-average-factor 7
        exit
#------------------------------------------
ALA-7>config>qos#
```

Editing QoS Policies

You can change existing policies and entries in the CLI or NMS. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors copy the policy to a work area, make the edits, and then write over the original policy.

Slope QoS Policy Command Reference

Command Hierarchies

Configuration Commands

```
config
    — qos
        — [no] slope-policy name
            — description description-string
            — no description
            — [no] high-slope
                — max-avg percent
                — no max-avg
                — max-prob percent
                — no max-prob
                — start-avg percent
                — no start-avg
                — [no] shutdown
            — [no] low-slope
                — max-avg percent
                — no max-avg
                — max-prob percent
                — no max-prob
                — start-avg percent
                — no start-avg
                — [no] shutdown
            — time-average-factor value
            — no time-average-factor
```

Operational Commands

```
config
    — qos
        — copy slope-policy src-name dst-name [overwrite]
```

Show Commands

show
— **qos**
— **slope-policy** [*slope-policy-name*] [**detail**]

Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
no description

Context config>qos>slope-policy

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax **copy slope-policy** *src-name dst-name* [**overwrite**]

Context config>qos

Description This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id.

The **copy** command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters **slope-policy —** Indicates that the source policy ID and the destination policy ID are slope policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

overwrite — Specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, an error will occur if the destination policy ID exists.

```
ALA-7>config>qos# copy slope-policy default sp1
MINOR: CLI Destination "sp1" exists - use {overwrite}.
ALA-7>config>qos#overwrite
```

Slope Policy QoS Commands

slope-policy

Syntax [**no**] **slope-policy** *name*

Context config>qos

Description This command enables the context to configure a QoS slope policy.

Default slope-policy “default”

Parameters *name* — The name of the slope policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

time-average-factor

Syntax **time-average-factor** *value*
no time-average-factor

Context config>qos>slope-policy

Description This command sets a weighting factor to calculate the new shared buffer average utilization after assigning buffers for a packet entering a queue. To derive the new shared buffer average utilization, the buffer pool takes a portion of the previous shared buffer average and adds it to the inverse portion of the instantaneous shared buffer utilization.

The **time-average-factor** command sets the weighting factor between the old shared buffer average utilization and the current shared buffer instantaneous utilization when calculating the new shared buffer average utilization

The TAF value applies to all high and low priority RED slopes for ingress and egress access buffer pools controlled by the slope policy.

The **no** form of this command restores the default setting.

Default **7** - Weighting instantaneous shared buffer utilization is 0.8%.

Parameters *value* — Represents the Time Average Factor (TAF), expressed as a decimal integer. The value specified for TAF affects the speed at which the shared buffer average utilization tracks the instantaneous shared buffer utilization. A low value weights the new shared buffer average utilization calculation more to the shared buffer instantaneous utilization, zero using it exclusively.

A high value weights the new shared buffer average utilization calculation more to the previous shared buffer average utilization value.

Values 0 — 15

Slope Policy QoS Policy Commands

high-slope

Syntax [**no**] **high-slope**

Context config>qos>slope-policy

Description The **high-slope** context contains the commands and parameters for defining the high priority Random Early Detection (RED) slope graph. Each buffer pool supports a high priority RED slope for managing access to the shared portion of the buffer pool for high priority or in-profile packets.

The **high-slope** parameters can be changed at any time and the affected buffer pool high priority RED slopes will be adjusted appropriately.

The **no** form of this command restores the high slope configuration commands to the default values. If the commands within **high-slope** are set to the default parameters, the **high-slope** node will not appear in save config and show config output unless the detail parameter is present.

low-slope

Syntax [**no**] **low-slope**

Context config>qos>slope-policy

Description The **low-slope** context contains the commands and parameters for defining the low priority Random Early Detection (RED) slope graph. Each buffer pool supports a low priority RED slope for managing access to the shared portion of the buffer pool for low priority or out-of-profile packets.

The **low-slope** parameters can be changed at any time and the affected buffer pool low priority RED slopes must be adjusted appropriately.

The **no** form of this command restores the low slope configuration commands to the default values. If the leaf commands within **low-slope** are set to the default parameters, the **low-slope** node will not appear in save config and show config output unless the detail parameter is present.

RED Slope Commands

max-avg

Syntax **max-avg** *percent*
no max-avg

Context config>qos>slope-policy>high-slope
config>qos>slope-policy>low-slope

Description Sets the low priority or high priority Random Early Detection (RED) slope position for the shared buffer average utilization value where the packet discard probability rises directly to one. The percent parameter is expressed as a percentage of the shared buffer size.

The **no** form of this command restores the max-avg value to the default setting. If the current **start-avg** setting is larger than the default, an error will occur and the max-avg setting will not be changed to the default.

Default **max-avg 90** — High slope default is 90% buffer utilization before discard probability is 1.
max-avg 75 — Low slope default is 75% buffer utilization before discard probability is 1.

Parameters *percent* — The percentage of the shared buffer space for the buffer pool at which point the drop probability becomes 1. The value entered must be greater or equal to the current setting of **start-avg**. If the entered value is smaller than the current value of **start-avg**, an error will occur and no change will take place.

Values 0 — 100

max-prob

Syntax **max-prob** *percent*
no max-prob

Context config>qos>slope-policy>high-slope
config>qos>slope-policy>low-slope

Description Sets the low priority or high priority Random Early Detection (RED) slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one. The percent parameter is expressed as a percentage of packet discard probability where always discard is a probability of 1. A **max-prob** value of 80 represents 80% of 1, or a packet discard probability of 0.8.

The **no** form of this command restores the **max-prob** value to the default setting.

Default **max-prob 80** — 80% maximum drop probability corresponding to the **max-avg.**

Parameters *percent* — The maximum drop probability percentage corresponding to the **max-avg,** expressed as a decimal integer.

Values 0 — 100

shutdown

Syntax [**no**] **shutdown**

Context config>qos>slope-policy>high-slope
config>qos>slope-policy>low-slope

Description This command enables or disables the administrative status of the Random Early Detection slope.

By default, all slopes are shutdown and have to be explicitly enabled (**no shutdown**).

The **no** form of this command administratively enables the RED slope.

Default **shutdown** - RED slope disabled implying a zero (0) drop probability

start-avg

Syntax **start-avg** *percent*
no start-avg

Context config>qos>slope-policy>high-slope
config>qos>slope-policy>low-slope

Description This command sets the low priority or high priority Random Early Detection (RED) slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero. The percent parameter is expressed as a percentage of the shared buffer size.

The **no** form of this command restores the start-avg value to the default setting. If the max-avg setting is smaller than the default, an error will occur and the start-avg setting will not be changed to the default.

queue

Syntax **queue** *queue-id* **drop-rate** *num*
no queue *queue-id*

Context config>qos>slope-policy>high-slope
config>qos>slope-policy>low-slope

Description Sets the low priority or high priority Random Early Detection (RED) slope drop-rate for the shared buffer per queue.

The **no** form of this command restores the drop-rate value to the default setting.

Default drop-rate 1 — High slope default is 1 (6.25 drop-rate) for all the queues, this implies that once the shared buffer utilization reaches the start-threshold level then packets egressing out from a particular queue would be dropped at 6.25% rate.

drop-rate 0 — Low slope default is 0 (100% drop-rate) for all the queues, this implies that once the shared buffer utilization reaches the start-threshold level then packets egressing out from a particular queue would be dropped at 100% rate.

Parameters *queue-id* — Specifies the ID of the queue for which the drop-rate is to be configured.

Values 1 — 8

drop-rate *num* **—** Specifies the drop rate to be configured.

Values 0 — 7

Show Commands

slope-policy

Syntax **slope-policy** [*slope-policy-name*] [**detail**]

Context show>qos

Description This command displays slope policy information.

Parameters *slope-policy-name* — The name of the slope policy.

detail — Displays detailed information about the slope policy.

Output **Slope QoS Policy Output Fields —** The following table describes slope QoS policy output fields.

Table 47: Show QoS Slope Policy Output Fields

Label	Description
`Policy`	The ID that uniquely identifies the policy.
`Description`	A string that identifies the policy's context in the configuration file.
`Time Avg`	The weighting between the previous shared buffer average utilization result and the new shared buffer utilization.
Slope Parameters	
`Start Avg`	Specifies the low priority or high priority RED slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero.
`Max Avg`	Specifies the percentage of the shared buffer space for the buffer pool at which point the drop probability becomes 1, expressed as a decimal integer
`Admin State`	`Up` – The administrative status of the RED slope is enabled. `Down` – The administrative status of the RED slope is disabled. Specifies the low priority or high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one.
`Max Prob.`	Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one.

Sample Output

```
A:C# show qos slope-policy 2
===============================================================================
QoS Slope Policy
===============================================================================
Policy         : 2
Time Avg       : 7
-------------------------------------------------------------------------------
High Slope Parameters
-------------------------------------------------------------------------------
Start Avg      : 70                            Admin State  : Enabled
Max Avg        : 90                            Max Prob.    : 100
-------------------------------------------------------------------------------
Low Slope Parameters
-------------------------------------------------------------------------------
Start Avg      : 30                            Admin State  : Enabled
Max Avg        : 40                            Max Prob.    : 100
===============================================================================

A:C# show qos slope-policy 2 detail
===============================================================================
QoS Slope Policy
===============================================================================
Policy         : 2
Time Avg       : 7
-------------------------------------------------------------------------------
High Slope Parameters
-------------------------------------------------------------------------------
Start Avg      : 70                            Admin State  : Enabled
Max Avg        : 90                            Max Prob.    : 100
-------------------------------------------------------------------------------
Low Slope Parameters
-------------------------------------------------------------------------------
Start Avg      : 30                            Admin State  : Enabled
Max Avg        : 40                            Max Prob.    : 100
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Object Type Object Id     Application          Pool
-------------------------------------------------------------------------------
Port        1/1/1         Acc-Egr              default
===============================================================================
A:C#
```

Shared-Queue QoS Policies

In This Section

This section provides information to configure shared-queue QoS policies using the command line interface.

Topics in this section include:

- Overview on page 630
- Basic Configurations on page 638
- Default Shared Queue Policy Values on page 643

Overview

Shared-queue QoS policies can be implemented to facilitate queue consumption on the router. It is especially useful when VPLS, IES, and VPRN services are scaled to very high numbers. Instead of allocating multiple hardware queues for each unicast queue defined in a SAP ingress QoS policy, SAPs with the shared-queuing feature enabled only allocate one hardware queue for each SAP ingress QoS policy unicast queue.

However, as a trade-off, the total amount of traffic throughput at the ingress of the node is reduced because any ingress packet serviced by a shared-queuing SAP is recirculated for further processing. This can reduce the bandwidth by half. Shared-queuing can add latency. Network planners should consider these restrictions while trying to scale services on the router.

Multipoint Shared Queuing

Multipoint shared queuing is supported to minimize the number of multipoint queues created for ingress VPLS, IES or VPRN SAPs or ingress subscriber SLA profiles. Normally, ingress multipoint packets are handled by multipoint queues created for each SAP or subscriber SLA profile instance. In some instances, the number of SAPs or SLA profile instances are sufficient for the in use multipoint queues to represent many thousands of queues on an ingress forwarding plane. If multipoint shared queuing is enabled for the SAPs or SLA profile instances on the forwarding plane, the multipoint queues are not created. Instead, the ingress multipoint packets are handled by the unicast queue mapped to the forwarding class of the multipoint packet.

Functionally, multipoint shared queuing is a superset of shared queuing. With shared queuing on a SAP or SLA profile instance, only unicast packets are processed twice, once for the initial service level queuing and a second time for switch fabric destination queuing. Shared queuing does not affect multipoint packet handling. Multipoint packet handling in normal (service queuing) is the same as shared queuing. When multipoint shared queuing is enabled, shared queuing for unicast packets is automatically enabled.

Ingress Queuing Modes of Operation

Three modes of ingress SAP queuing are supported for multipoint services (IES, VPLS and VPRN); service, shared, and multipoint shared. The same ingress queuing options are available for IES and VPLS subscriber SLA profile instance queuing.

Ingress Service Queuing

Normal or service queuing is the default mode of operation for SAP ingress queuing. Service queuing preserves ingress forwarding bandwidth by allowing a service queue defined in an ingress SAP QoS policy to be represented by a group of hardware queues. A hardware queue is created for each switch fabric destination to which the logical service queue must forward packets. For a VPLS SAP with two ingress unicast service queues, two hardware queues are used for each destination forwarding engine the VPLS SAP is forwarding to. If three switch fabric destinations are involved, six queues are allocated (2 unicast service queues multiplied by 3 destination forwarding complexes equals six hardware queues). Figure 25 demonstrates unicast hardware queue expansion. Service multipoint queues in the ingress SAP QoS policy are not expanded to multiple hardware queues, each service multipoint queue defined on the SAP equates to a single hardware queue.

When multiple hardware queues represent a single logical service queue, the system automatically monitors the offered load and forwarding rate of each hardware queue. Based on the monitored state of each hardware queue, the system imposes an individual CIR and PIR rate for each queue that provides an overall aggregate CIR and PIR reflective of what is provisioned on the service queue.

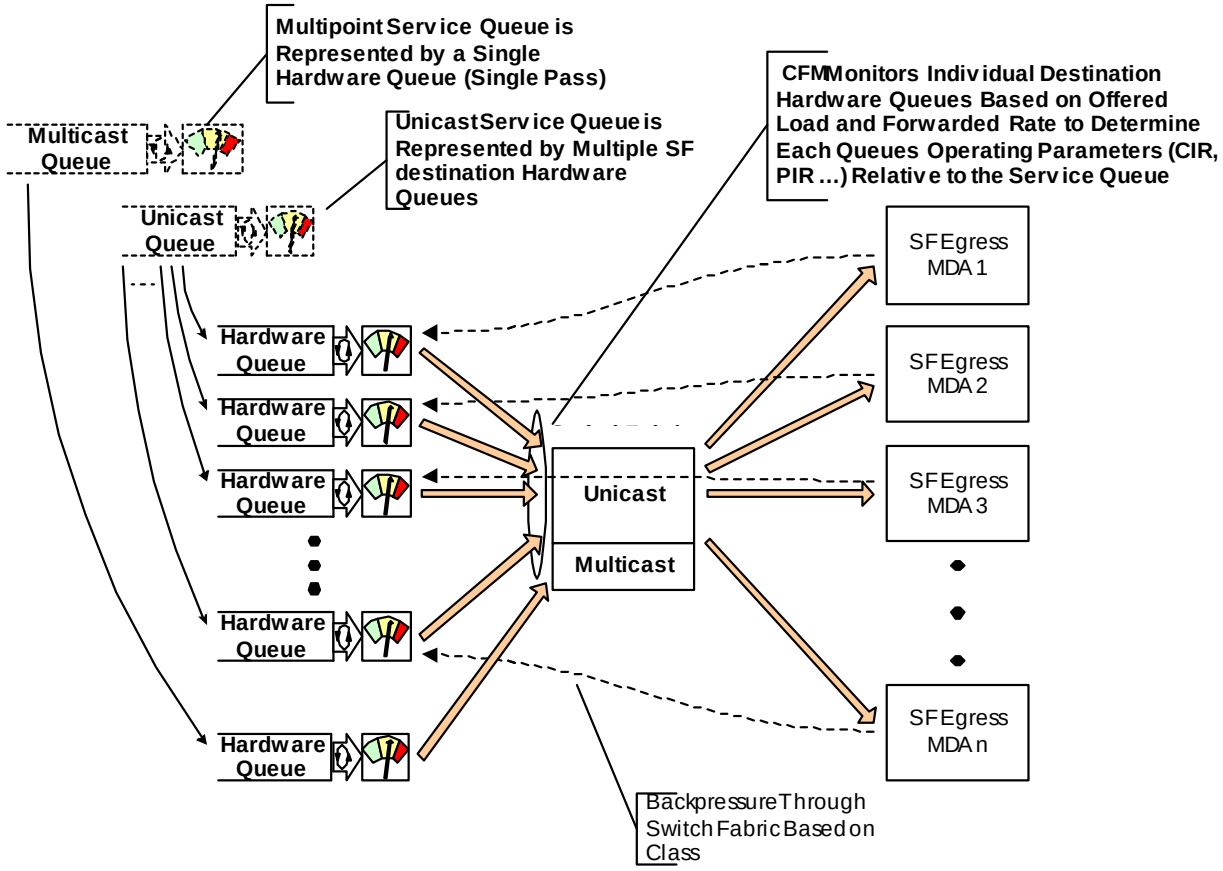

Figure 25: Unicast Service Queue Mapping to Multiple Destination Based Hardware Queues

Ingress Shared Queuing

To avoid the hardware queue expansion issues associated with normal service based queuing, the system allows an ingress logical service queue to map to a single hardware queue when shared queuing is enabled. Shared queuing uses two passes through the ingress forwarding plane to separate ingress per service queuing from the destination switch fabric queuing. In the case of shared queuing, ingress unicast service queues are created one-for-one relative to hardware queues. Each hardware queue representing a service queue is mapped to a special destination in the traffic manager that 'forwards' the packet back to the ingress forwarding plane allowing a second pass through the traffic manager. In the second pass, the packet is placed into a 'shared' queue for the destination forwarding plane. The shared queues are used by all services configured for shared queuing.

When the first SAP or SLA profile instance is configured for shared queuing on an ingress forwarding plane, the system allocates eight hardware queues per available destination forwarding plane, one queue per forwarding class. Twenty-four hardware queues are also allocated for multipoint shared traffic, but that is discussed in the following section. The shared queue parameters that define the relative operation of the forwarding class queues are derived from the Shared Queue policy defined in the QoS CLI node. Figure 26 demonstrates shared unicast queuing. SAP or SLA profile instance multipoint queuing is not affected by enabling shared queuing. Multipoint queues are still created as defined in the ingress SAP QoS policy and ingress multipoint packets only traverse the ingress forwarding plane a single time.

Enabling shared queuing may affect ingress performance due to double packet processing through the service and shared queues.

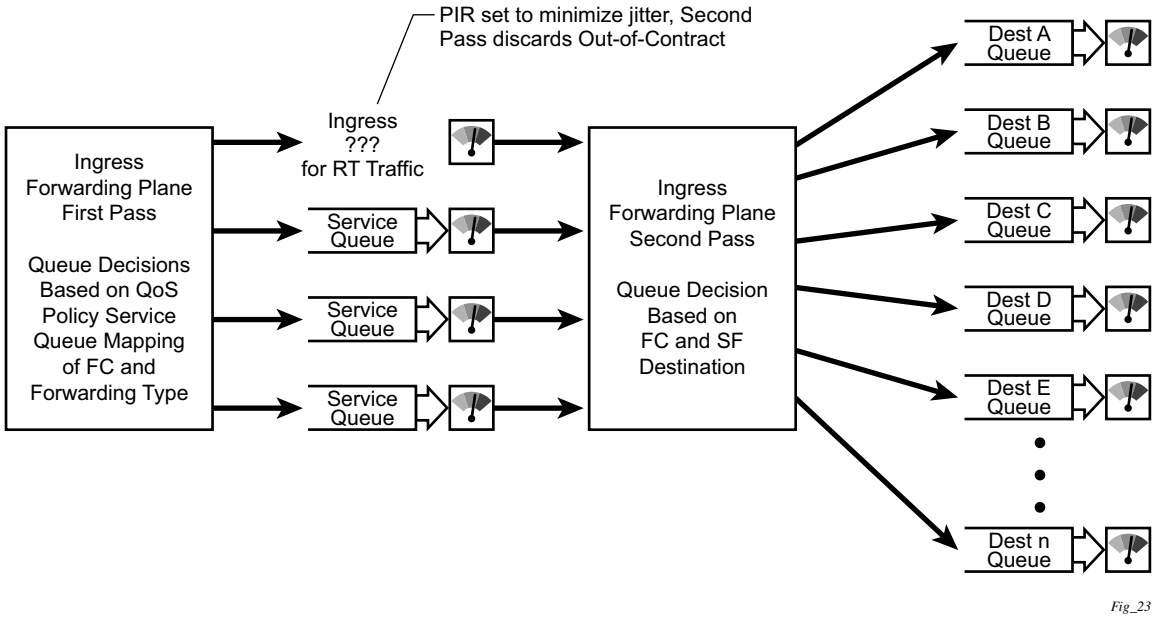

Figure 26: Unicast Service Queuing With Shared Queuing Enabled

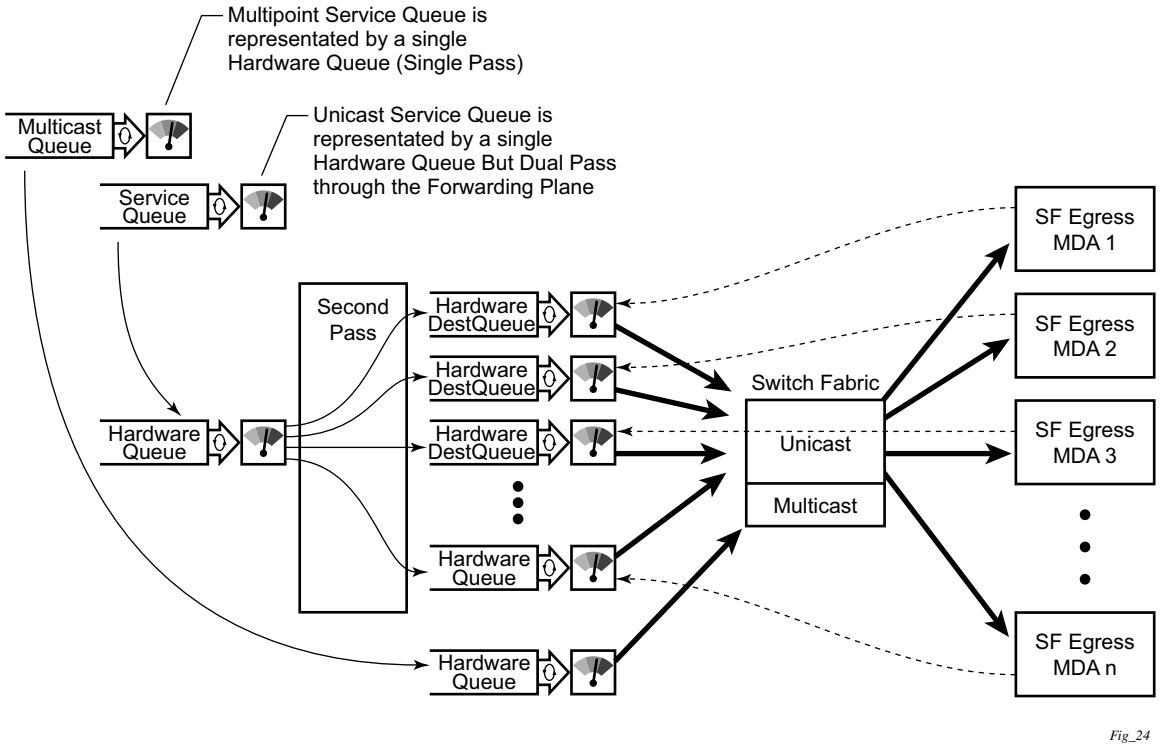

Figure 27: Multipoint Queue Behavior with Shared Queuing Enabled

Ingress Multipoint Shared Queuing

Ingress multipoint shared queuing is a variation to the unicast shared queuing defined in Ingress Shared Queuing on page 633. Ingress unicast service queues are mapped one-for-one with hardware queues and unicast packets traverse the ingress forwarding plane twice. In addition to the above, the multipoint queues defined in the ingress SAP QoS policy are not created. Instead, multipoint packets (broadcast, multicast and unknown unicast destined) are treated to the same dual pass ingress forwarding plane processing as unicast packets. In the first pass, the forwarding plane uses the unicast queue mappings for each forwarding plane. The second pass uses the multipoint shared queues to forward the packet to the switch fabric for special replication to all egress forwarding planes that need to process the packet.

The benefit of defining multipoint shared queuing is the savings of the multipoint queues per service. By using the unicast queues in the first pass and then the aggregate shared queues in the second pass, per service multipoint queues are not required. The predominant scenario where multipoint shared queuing may be required is with subscriber managed QoS environments using a subscriber per SAP model. Usually, ingress multipoint traffic is minimal per subscriber and the extra multipoint queues for each subscriber reduces the overall subscriber density on the ingress forwarding plane. Multipoint shared queuing eliminates the multipoint queues sparing hardware queues for better subscriber density. Figure 2.3 demonstrates multipoint shared queuing.

One caveat of enabling multipoint shared queuing is that multipoint packets are no longer managed per service (although the unicast forwarding queues may provide limit benefit in this area). Multipoint packets in a multipoint service (VPLS, IES and VPRN) use significant resources in the system, consuming ingress forwarding plane multicast bandwidth and egress replication bandwidth. Usually, the per service unicast forwarding queues are not rate limited to a degree that allows adequate management of multipoint packets traversing them when multipoint shared queuing is enabled. It is possible to minimize the amount of aggregate multipoint bandwidth by setting restrictions on the multipoint queue parameters in the QoS nodes Shared Queue policy. Aggregate multipoint traffic can be managed per forwarding class for each of the three forwarding types (broadcast, multicast or unknown unicast – broadcast and unknown unicast are only used by VPLS).

Another caveat for multipoint shared queuing is that multipoint traffic now consumes double the ingress forwarding plane bandwidth due to dual pass ingress processing.

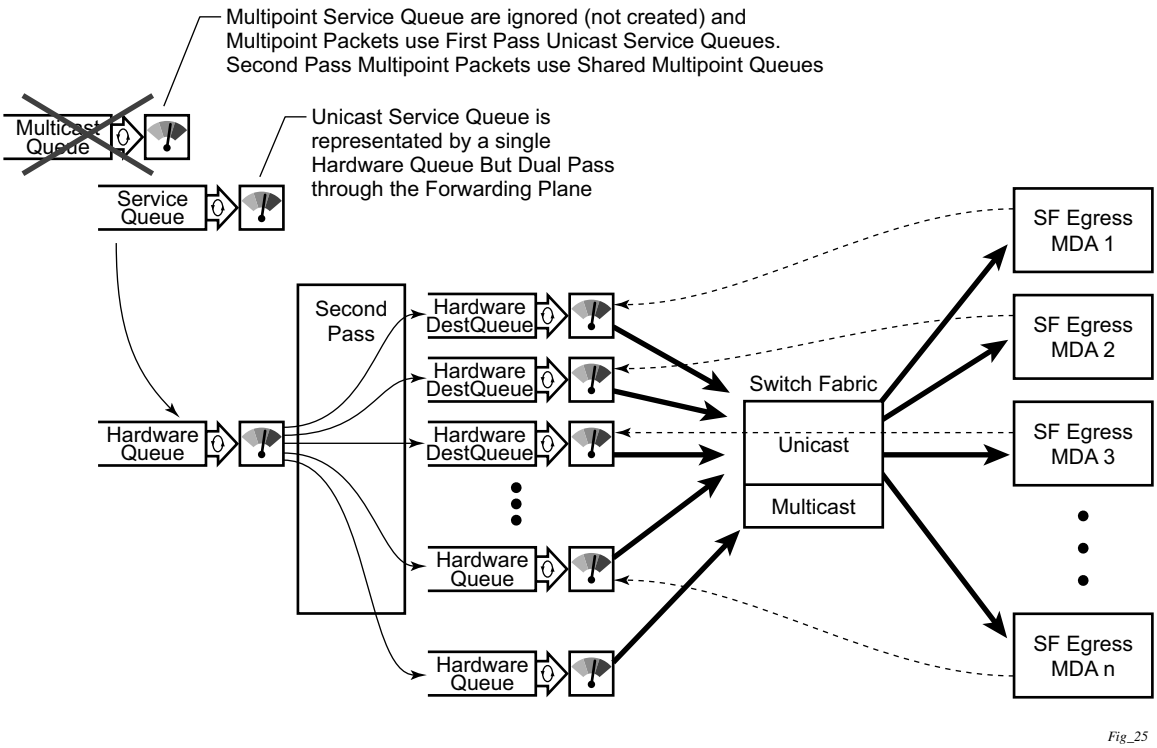

Figure 28: Multipoint Shared Queuing Using First Pass Unicast Queues

Note that multipoint shared queuing cannot be enabled on the following services:

- Epipe
- Apipe
- Fpipe
- Ipipe
- Routed CO

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your router, refer to CLI Usage chapter in the Basic System Configuration Guide.

Basic Configurations

The default shared queue QoS policy conforms to the following:

- There is only one default shared queue policy in the system.
- The default shared queue policy has fixed forwarding classes, queues and FC-queue mapping, which cannot be modified, added, or deleted.

The only configurable entities in the default shared queue policy are the queue attributes, queue priority, and the description string. The queue priority for a shared queue can be changed to expedited, best-effort or auto-expedited.

Modifying the Default Shared-Queue Policy

The only configurable entities in the default shared queue policy are the queue attributes and the description string. The changes are applied immediately to all services where this policy is applied. Use the following CLI syntax to modify a shared-queue policy:

CLI Syntax:

```
config>qos#
   shared-queue name
      description description-string
      queue queue-id [queue-type] [multipoint]
         cbs percent
         high-prio-only percent
         mbs percent
         rate percent [cir percent]
```

The following displays a shared-queue policy configuration example:

```
A:ALA-48>config>qos>shared-queue# info
----------------------------------------------
            description "test1"
            queue 1 create
                cbs 2
                high-prio-only 20
            exit
----------------------------------------------
A:ALA-48>config>qos>shared-queue#
```

Applying Shared-Queue Policies

The default shared queue policy is applied at the SAP level just as sap-ingress and sap-egress QoS policies are specified. If the **shared-queuing** keyword is not specified in the **qos** *policy-id* command then the SAP is assumed to use single-pass queuing.

Apply shared-queue policies to the following entities:

- Epipe Services
- IES Services
- VPLS Services
- VPRN Services

Epipe Services

Use the following CLI syntax to apply QoS policies to ingress Epipe SAPs:

CLI Syntax:
```
config>service>epipe service-id [customer customer-id]
    sap sap-id
        ingress
            qos policy-id [shared-queuing]
```

The following output displays an Epipe service configuration with SAP ingress policy 100 applied to the SAP with shared-queuing enabled.

```
A:SR>config>service# info
----------------------------------------------
  epipe 6 customer 6 vpn 6 create
     description "Distributed Epipe to west coast"
     sap 1/1/10:0 create
          ingress
               qos 100 shared-queuing
          exit
     exit
     no shutdown
  exit
----------------------------------------------
A:SR>config>service#
```

IES Services

Use the following CLI syntax to apply the default policy to an IES service:

CLI Syntax:
```
config>service# ies service-id
             interface interface-name
                      sap sap-id
                         ingress
                            qos policy-id [shared-queuing |multipoint-
                               shared]
```

The following output displays an IES service configuration with SAP ingress policy 100 applied to the SAP with shared-queuing enabled.

```
A:SR>config>service# info
----------------------------------------------
   ies 88 customer 8 vpn 88 create
     interface "Sector A" create
          sap 1/1/1.2.2 create
               ingress
                    qos 100 multipoint-shared
               exit
          exit
     exit
     no shutdown
   exit
----------------------------------------------
A:SR>config>service#
```

VPLS Services

Use the following CLI syntax to apply the default shared-queue policy to an ingress VPLS SAP:

CLI Syntax:
```
config>service# vpls service-id [customer customer-id]
                 sap sap-id
                    ingress
                       qos policy-id [shared-queuing | multipoint-shared]
```

The following output displays a VPLS service configuration with SAP ingress policy 100 with shared-queuing enabled.

```
A:SR>config>service# info
----------------------------------------------
  vpls 700 customer 7 vpn 700 create
     description "test"
     sap 1/1/9:0 create
          ingress
               qos 100 multipoint-shared
          exit
     exit
  exit
----------------------------------------------
A:SR>config>service#
```

VPRN Services

Use the following CLI syntax to apply QoS policies to ingress VPRN SAPs:

CLI Syntax:
```
config>service# vprn service-id [customer customer-id]
                 interface ip-int-name
                    sap sap-id
                       ingress
                          qos policy-id [shared-queuing | multipoint-shared]
```

The following output displays a VPRN service configuration. The default SAP ingress policy was not modified but shared queuing was enabled.

```
A:SR7>config>service# info
----------------------------------------------
  vprn 1 customer 1 create
     interface "to-ce1" create
          address 11.1.0.1/24
          sap 1/1/10:1 create
               ingress
                    qos 1 multipoint-shared
               exit
```

```
            exit
        exit
        no shutdown
    exit
-------------------------------------------------
A:SR7>config>service#
```

Default Shared Queue Policy Values

The only allowed shared queue policy is the default and cannot be deleted. The only configurable entities are the queue priority, attributes of individual queues and the description string. Table 48 lists the default values.

Table 48: Shared Queue Policy Defaults

Field	Default
description	“Default Shared Queue Policy”
queue 1	auto-expedite
rate	100
cir	0
mbs	50
cbs	1
high-prio-only	10
queue 2	auto-expedite
rate	100
cir	25
mbs	50
cbs	3
high-prio-only	10
queue 3	auto-expedite
rate	100
cir	25
mbs	50
cbs	10
high-prio-only	10
queue 4	auto-expedite
rate	100
cir	25
mbs	25
cbs	3
high-prio-only	10
queue 5	auto-expedite
rate	100
cir	100
mbs	50
cbs	10

Table 48: Shared Queue Policy Defaults (Continued)

Field	Default
high-prio-only	10
queue 6	auto-expedite
rate	100
cir	100
mbs	50
cbs	10
high-prio-only	10
queue 7	auto-expedite
rate	100
cir	10
mbs	25
cbs	3
high-prio-only	10
queue 8	auto-expedite
rate	100
cir	10
mbs	25
cbs	3
high-prio-only	10

The fc-to-shared-queue mappings that cannot be modified are:

fc af	queue 3
fc be	queue 1
fc h1	queue 6
fc h2	queue 5
fc l1	queue 4
fc l2	queue 2
fc nc	queue 8

The following output displays the default configuration:

```
ALA-7>config>qos>shared-queue# info detail
----------------------------------------------
            description "Default Shared Queue Policy"
            queue 1 auto-expedite create
                rate 100 cir 0
                mbs 50
                cbs 1
                high-prio-only 10
            exit
            queue 2 auto-expedite create
                rate 100 cir 25
                mbs 50
                cbs 3
                high-prio-only 10
            exit
            queue 3 auto-expedite create
                rate 100 cir 25
                mbs 50
                cbs 10
                high-prio-only 10
            exit
            queue 4 auto-expedite create
                rate 100 cir 25
                mbs 25
                cbs 3
                high-prio-only 10
            exit
            queue 5 auto-expedite create
                rate 100 cir 100
                mbs 50
                cbs 10
                high-prio-only 10
            exit
            queue 6 auto-expedite create
                rate 100 cir 100
                mbs 50
                cbs 10
                high-prio-only 10
            exit
            queue 7 auto-expedite create
                rate 100 cir 10
                mbs 25
                cbs 3
                high-prio-only 10
            exit
            queue 8 auto-expedite create
                rate 100 cir 10
                mbs 25
                cbs 3
                high-prio-only 10
            exit
            fc af create
                queue 3
            exit
            fc be create
                queue 1
```

```
            exit
            fc ef create
                queue 6
            exit
            fc h1 create
                queue 7
            exit
            fc h2 create
                queue 5
            exit
            fc l1 create
                queue 4
            exit
            fc l2 create
                queue 2
            exit
            fc nc create
                queue 8
            exit
-------------------------------------------------------------------------------
ALA-7>config>qos>shared-queue#
```

Shared-Queue QoS Policy Command Reference

Command Hierarchies

Configuration Commands

```
config
    — qos
        — shared-queue policy-name
            — [no] clp-tagging
            — description description-string
            — no description
            — [no] fc {be | l2 | af | l1 | h2 | ef | h1 | nc}
                — broadcast-queue queue-id
                — no broadcast-queue
                — multicast-queue queue-id
                — no multicast-queue
                — queue queue-id
                — no queue
                — unknown-queue queue-id
                — no unknown-queue
            — queue queue-id [queue-type] [profile-mode | priority-mode] [multipoint]
            — no queue queue-id
                — cbs percent
                — no cbs
                — high-prio-only percent
                — no high-prio-only
                — mbs percent
                — no mbs
                — rate percent [cir percent]
                — no rate
```

Show Commands

```
show
    — qos
        — shared-queue [policy-name] [detail]
```

Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
no description

Context config>qos>shared-queue

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Shared Queue QoS Commands

shared-queue

Syntax **shared-queue** *policy-name*

Context config>qos

Description This command enables the context to modify the QoS **default** shared-queue policy.

Parameters *policy-name* — The name of the **default** shared-queue policy.

Values **default**

fc

Syntax [**no**] **fc** {**be** | **l2** | **af** | **l1** | **h2** | **ef** | **h1** | **nc**}

Context config>qos>shared-queue

Description This command specifies the forwarding class name. The forwarding class name represents an egress queue. The **fc** *fc-name* represents a CLI parent node that contains sub-commands or parameters describing the egress characteristics of the queue and the marking criteria of packets flowing through it. The **fc** command overrides the default parameters for that forwarding class defined in the network default policy *policy-id* 1.

Default See Default Shared Queue Policy Values on page 643 for undefined forwarding class values.

Parameters *fc-name* — The case-sensitive, system-defined forwarding class name for which policy entries will be created.

Default none

broadcast-queue

Syntax **broadcast-queue** *multipoint-queue-id*

Context config>qos>shared-queue>fc

Description This command configures the broadcast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all broadcast traffic on a SAP using this policy will be forwarded using the *queue-id*.

The broadcast forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of the command sets the broadcast forwarding type *queue-id* back to the default of tracking the multicast forwarding type queue mapping.

Parameters *queue-id* — The *queue-id* parameter must be an existing, multipoint queue defined in the **config>qos>sap-ingress** context .

Values 17 — 24

multicast-queue

Syntax **multicast-queue** *queue-id*

Context config>qos>shared-queue>fc

Description This command configures the multicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all multicast traffic on a SAP using this policy is forwarded using the *queue-id*.

The multicast forwarding type includes the **unknown** unicast forwarding type and the **broadcast** forwarding type unless each is explicitly defined to a different multipoint queue. When the unknown and broadcast forwarding types are left as default, they will track the defined queue for the multicast forwarding type.

The **no** form of the command sets the multicast forwarding type *queue-id* back to the default queue for the forwarding class. If the **broadcast** and **unknown** forwarding types were not explicitly defined to a multipoint queue, they will also be set back to the default multipoint queue (queue 11).

Parameters *queue-id* — The *queue-id* parameter specified must be an existing, multipoint queue defined in the the **config>qos>sap-ingress** context.

Values 9 — 16

Default 11

queue

Syntax **queue** *queue-id*
no queue

Context config>qos>shared-queue>fc

This command overrides the default unicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a non-multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all unicast traffic (this includes all traffic, even broadcast and multicast for services) on a SAP using this policy is forwarded using the *queue-id*.

The **no** form of this command sets the unicast (point-to-point) *queue-id* back to the default queue for the forwarding class (queue 1).

Parameters *queue-id* — The *queue-id* parameter specified must be an existing, non-multipoint queue defined in the **config>qos>sap-ingress** context.

Values Any valid non-multipoint *queue-id* in the policy including 1 and 3 through 32 .

Default 1

queue

Syntax **queue** *queue-id* [*queue-type*] [**profile-mode | priority-mode**] [**multipoint**]
queue *queue-id* [*queue-type*] [**multipoint**]
no queue *queue-id*

Context config>qos>shared-queue

Description This command creates the context to configure a shared queue QoS policy queue.
Explicit definition of an ingress queue's hardware scheduler status is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or non-expedited nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (`nc`, `ef`, `h1` or `h2`), the queue is treated as an expedited queue by the hardware schedulers. When any non-expedited forwarding classes are mapped to the queue (`be`, `af`, `l1` or `l2`), the queue is treated as best effort (`be`) by the hardware schedulers. The hardware status of the queue must be defined at the time of queue creation within the policy.

Parameters *queue-id* — The *queue-id* for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

Values 1 — 32

queue-type — The **expedite**, **best-effort** and **auto-expedite** queue types are mutually exclusive to each other. Each defines the method that the system uses to service the queue from a hardware perspective. While parental virtual schedulers can be defined for the queue, they only enforce how the queue interacts for bandwidth with other queues associated with the same scheduler hierarchy. An internal mechanism that provides access rules when the queue is vying for bandwidth with queues in other virtual schedulers is also needed. A keyword must be specified at the time the queue is created in the SAP ingress policy. If an attempt to change the keyword after the queue is initially defined, an error is generated.

expedite — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

auto-expedite — This keyword allows the system to auto-define the way the queue is serviced by the hardware. When **auto-expedite** is defined on the queue, the queue is treated in an expedited manner when all forwarding classes mapped to the queue are configured as expedited types `nc`, `ef`, `h1` or `h2`. When a single non-expedited forwarding class is mapped to the queue (`be`, `af`, `l1` and `l2`) the queue automatically falls back to non-expedited status.

Default auto-expedite

multipoint — This keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. If you attempt to map forwarding class unicast traffic to a multipoint queue, an error is generated and no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a pre-existing multipoint queue to edit *queue-id* parameters.

unknown-queue

Syntax **unknown-queue** *queue-id*
no unknown-queue

Context config>qos>shared-queue>fc

Description This command configures the unknown unicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the *queue-id*.

The unknown forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of this command sets the unknown forwarding type *queue-id* back to the default of tracking the multicast forwarding type queue mapping.

Parameters *queue-id* — The *queue-id* must be an existing, multipoint queue defined in the the **config>qos>sap-ingress** context .

Values 25 — 32 .

cbs

Syntax **cbs** *percent*
no cbs

Context config>qos>shared-queue>queue

Description The Committed Burst Size (**cbs**) command specifies the relative amount of reserved buffers for a specific ingress network MDA forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.

The CBS for a queue is used to determine whether it has exhausted its reserved buffers while enqueuing packets. Once the queue has exceeded the amount of buffers considered in reserve for this queue, it must contend with other queues for the available shared buffer space within the buffer pool. Access to this shared pool space is controlled through Random Early Detection (RED) slope application.

Two RED slopes are maintained in each buffer pool. A high priority slope is used by in-profile packets. A low priority slope is used by out-of-profile packets. All Network-Control and Management packets are considered in-profile. Assured packets are handled by their in-profile and out-of-profile markings. All Best-Effort packets are considered out-of-profile. Premium queues should be configured such that the CBS percent is sufficient to prevent shared buffering of packets. This is generally taken care of by the CIR scheduling of Premium queues and the overall small amount of traffic on the class. Premium queues in a properly designed system will drain before all others, limiting their buffer utilization.

The RED slopes will detect congestion conditions and work to discard packets and slow down random TCP session flows through the queue. The RED slope definitions can be defined, modified or

disabled through the network-queue policy assigned to the MDA for the network ingress buffer pool or assigned to the network port for network egress buffer pools.

The resultant CBS size can be larger than the MBS. This will result in a portion of the CBS for the queue to be unused and should be avoided.

Default The **cbs** forwarding class defaults are listed in the table below:

Forwarding Class	Fowarding Class Label	Default CBS
Network-Control	nc	3
High-1	h1	3
Expedited	ef	1
High-2	h2	1
Low-1	l1	3
Assured	af	1
Low-2	l2	3
Best-Effort	be	1

Parameters *percent* — The percent of buffers reserved from the total buffer pool space, expressed as a decimal integer. If 10 MB is the total buffers in the buffer pool, a value of 10 would reserve 1MB (10%) of buffer space for the forwarding class queue. The value 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can be applied for scheduling purposes).

Values 0 — 100

high-prio-only

Syntax **high-prio-only** *percent*
no high-prio-only

Context config>qos>shared-queue>queue

Description The **high-prio-only** command allows the reservation of queue buffers for use exclusively by high priority packets as a default condition for access buffer queues for this shared queue policy.

The difference between the MBS size for the queue and the high priority reserve defines the threshold where low priority traffic will be discarded. The result is used on the queue to define a threshold where low priority packets are discarded, leaving the rest of the default MBS size for high priority packets only. If the current MBS for the queue is 10MBytes, a value of 5 will result in a high priority reserve on the queue of 500KBytes. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.

Modifying the current MBS for the queue through the **mbs** command will cause the default **high-prio-only** function to be recalculated and applied to the queue. The **high-prio-only** command as defined for the specific queue can be used to override the default **high-prio-only** setting as defined in

the network queue policy. This prevents the **high-prio-only** command for the shared queue policy from having an affect on the queue.

Default The **high-prio-only** forwarding class defaults are listed in the table below.

Forwarding Class	Fowarding Class Label	Default high-prio-only
Network-Control	nc	10
High-1	h1	10
Expedited	ef	10
High-2	h2	10
Low-1	l1	10
Assured	af	10
Low-2	l2	10
Best-Effort	be	10

Parameters *percent* — The amount of queue buffer space, expressed as a decimal percentage of the MBS.

Values 0 — 100 | default

mbs

Syntax **mbs** *percent*
no mbs

Context config>qos>shared-queue>queue

Description This command specifies the relative amount of the buffer pool space for the maximum buffers for a specific ingress network MDA forwarding class queue or egress network port forwarding class queue.

The MBS value is used to by a queue to determine whether it has exhausted its total allowed buffers while enqueuing packets. Once the queue has exceeded its maximum amount of buffers, all packets are discarded until the queue transmits a packet. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of the network queues.

The MBS size can sometimes be smaller than the CBS. This will result in a portion of the CBS for the queue to be unused and should be avoided.

Default The **mbs** forwarding class defaults are listed in the table below.

Forwarding Class	Fowarding Class Label	Default MBS
Network-Control	nc	25
High-1	h1	25
Expedited	ef	50
High-2	h2	50
Low-1	l1	25
Assured	af	50
Low-2	l2	50
Best-Effort	be	50

Parameters *percent* — The percent of buffers from the total buffer pool space for the maximum amount of buffers, expressed as a decimal integer. If 10 MB is the total buffers in the buffer pool, a value of 10 would limit the maximum queue size to 1MB (10%) of buffer space for the forwarding class queue. If the total size is increased to 20MB, the existing value of 10 would automatically increase the maximum size of the queue to 2MB.

Values 0 — 100

rate

Syntax **rate** [*percent*] [**cir** *percent*]
no rate

Context config>qos>shared-queue>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the percentage that the queue can transmit packets (for SAP ingress queues) or out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the percentage at which the system prioritizes the queue over other queues competing for the same bandwidth. For SAP ingress, the CIR also defines the rate that packets are considered in-profile by the system. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP ingress or SAP egress QoS policy with the *queue-id*.

Parameters *percent* — Defines the percentage of the max rate allowed for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been

executed, the default PIR of **max** is assumed. Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 0 — 100, **max**

Default **100**

cir *percent* **—** Defines the percentage of the max rate allowed for the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer.

Values 0 — 100, **max**

Default 0

Show Commands

shared-queue

Syntax **shared-queue** [*shared-queue-policy-name*] [**detail**]

Context show>qos

Description This command displays shared-queue policy information.

Parameters *shared-queue-policy-name* — The shared-queue policy name.

detail — Displays detailed information about the shared-queue policy.

Output **Shared-Queue QoS Policy Output Fields —** The following table describes shared-queue QoS policy output fields.

Table 49: Show QoS Shared Queue Output Fields

Label	Description
Policy	The ID that uniquely identifies the policy.
Description	A text string that helps identify the policy's context in the configuration file.

Sample Output

```
A:ALA-1>config>qos# show qos shared-queue default
===============================================================================
QoS Network Queue Policy
===============================================================================
-------------------------------------------------------------------------------
Shared Queue Policy (default)
-------------------------------------------------------------------------------
Policy         : default
Description    : Default Shared Queue Policy
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
A:ALA-1>config>qos#
```

QoS ATM Traffic Descriptor Profiles

In This Section

This section provides information to configure QoS Traffic Descriptor Profiles using the command line interface.

- Overview on page 660
- Basic Configurations on page 666
- Default ATM-TD-Profile Policy Values on page 669
- Service Management Tasks on page 670

Overview

ATM Traffic Descriptor Profiles

This section provides a description of support ATM QoS policy features. Each traffic descriptor defines the expected rates and characteristics of traffic.

ATM Traffic Management

The router supports the ATM Forum Traffic Management Specification Version 4.1. The following sections describe the QoS features for ATM Permanent Virtual Connection (PVC).

QoS Model for ATM-Based Services

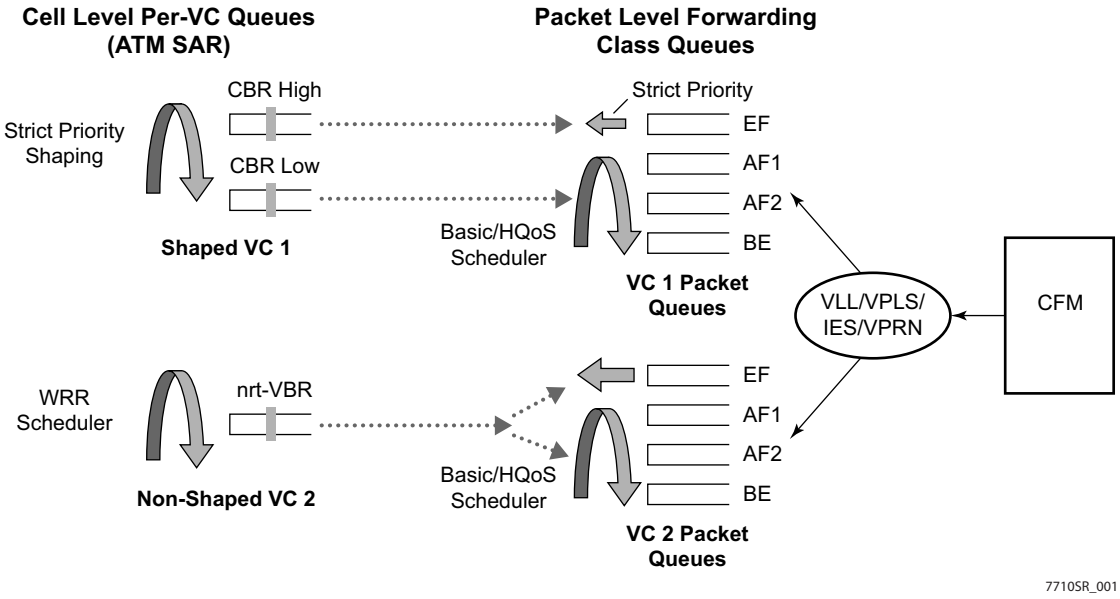

Figure 29: Hierarchical Scheduling for ATM-Based Services

This section provides a description of the QoS model used for ATM-based services on the router ATM MDA. Although slight variations of this model are applied on other ATM capable MDAs, the principles remain the same. An example of a VPLS Service with ATM SAP is shown in Figure 29.

When Ethernet frames are sent over an ATM VC, the scheduling of data becomes hierarchical with two main levels: packet level scheduling and per-VC cell level scheduling.

At the first level, frames are queued on a per-CoS (or forwarding class), per-VC basis in order to achieve the proper class of service differentiation for the frames in the same VC. Each Ethernet frame is queued based on the VC dedicated forwarding class queue, as configured in the service egress QoS policy. The packet level scheduling can make use of HQoS scheduler policy in order to enforce aggregate bandwidth among a group of queues feeding an ATM VC or to enforce aggregation of bandwidth across all queues of all VCs at a given customer site.

The frame level scheduling is the same for other types of SAP (Ethernet, FR, and PPP) and all the features available on the service ingress and egress QoS policies can be applied.

At the second level, the segmented cells are queued in per-VC queues according to the service category of the ATM traffic descriptor profile applied to the ATM SAP. Scheduling at the ATM level enforces the priority and bandwidth sharing desired at the cell level.

It is important to note that any discard decision are performed exclusively at the packet level where the context for the frame forwarding class and for the 802.1p bit mapping to a forwarding class is known. When a per-VC queue backs up, a back pressure scheme is applied such that the frames are held in the per-forwarding class packet queues dedicated to this VC.

This hierarchical scheduling of frames and cells of a given VC terminating on a VPLS instance provides the flexibility to apply policing and shaping on a per-forwarding class basis for the Ethernet frames of each VC as well as the option to shape the aggregate cell flow into the ATM VC back into the customer site.

This QoS model is applied to all router services which include an ATM SAP (VLL, VPLS, IES and VPRN).

ATM Service Categories

The router supports the following service categories,

- CBR - Constant Bit Rate
- Rt-VBR - Real-Time Variable Bit Rate
- nrt-VBR - Non Real-Time Variable Bit Rate
- UBR/UBR+MIR - Unspecified Bit Rate with Minimum Cell Rate. Note that UBR is a special case of UBR+MIR where MIR=0.

ATM Traffic Descriptors and QoS Parameters

The router supports the following traffic descriptors for ATM.

Table 50: ATM Traffic Descriptors

Service Category	Traffic Descriptors
CBR	P0_1 PIR in Kbps (applies to CLP=0 & CLP=1 flows)
Rt-VBR and nrt-VBR	P0_1and S0_1 PIR in Kbps (applies to CLP=0 & CLP=1 flows) SIR in Kbps (applies to CLP=0 & CLP=1 flows) MBS in cells (applies to CLP=0 and CLP=1 flows) P0_1 and S0 PIR in Kbps (applies to CLP=0 & CLP=1 flows) SIR in Kbps (applies to CLP=0 flow only) MBS in cells (applies to CLP=0 flow only)
UBR/UBR+MIR	P0_1 PIR in Kbps (applies to CLP=0 & CLP=1 flows) MIR in Kbps (applies to CLP=0 & CLP=1 flows)

Policing

The policing option, when enabled, applies only for ingress traffic. Similarly, the shaping option, if enabled, applies only for egress traffic. For example, if a traffic descriptor has both options, policing and shaping enabled, the policing option is enforced for the ingress traffic, while the shaping option is enforced for the egress traffic. The policing option is valid for all service categories. The following ATM service category conformance definitions are supported:

- P0_1 - CBR, UBR
- P0_1andS0_1 – VBR.1
- P0_1andS0 – VBR.2
- P0_1andS0_Tag – VBR.3

Shaping

- Ingress shaping — ATM layer ingress shaping is not supported. Packet level shaping is supported as per the service ingress QoS policy applied to the ATM SAP.
- Egress shaping — ATM layer egress shaping is supported for CBR, rt-VBR, and nrt-VBR VCs. A CBR VC is shaped to a single leaky bucket with parameter PIR. A rt-VBR VC or a nrt-VBR VC is shaped to two leaky buckets with parameters PIR and {SIR, BT}, where

BT is the Burst Tolerance and is a function of the MBS parameters configured by the user in the traffic descriptor.

In the egress direction, packet level shaping is supported as per the service egress QoS policy applied to the ATM SAP.

ATM Queuing and Scheduling

The router provides a per-VC queuing architecture in the ATM-capable MDAs. In the egress direction towards the ATM port, the scheduling priority at the ATM layer is as follows:

- CBR VCs are scheduled with strict priority over all other service categories
- rt-VBR VCs are scheduled next with strict priority over nrt-VBR and UBR VCs.
- nrt-VBR shaped VCs are scheduled next with strict priority over nrt-VBR unshaped VCs and UBR VCs.
- nrt-VBR unshaped VCs and UBR VCs are scheduled as a common class. Scheduling among these VCs is done using a WRR scheduler where the weight of each VC is determined by the configured SIR for nrt-VBR and by the MIR for UBR VCs. The scheduling is work-conserving, so each VC has access to excess bandwidth in proportion to its SIR/MIR. Under congestion, the performance of each VC degrades proportional to the weight of the VC.

Congestion Avoidance

- PPD — An ATM cell discarded in the middle of an AAL5 packet makes the entire packet unusable. The PPD mechanism attempts to minimize the congestion problems that can occur at the higher layers (TCP) due to ATM cell discards.

  With PPD, once a cell is discarded by the ATM policing function, no other cells for that PDU are accepted, with exception of the “tail cell” which is sent to inform the far end that the end of a frame has arrived.

  PPD is enabled on ATM SAP part of an AAL5 SDU Apipe VLL service and on all services where the ATM SAP cell stream is re-assembled, i.e., IES, VPRN, VPLS, Epipe, and Apipe services.
- WRED — Congestion and potential discards are performed on per forwarding class basis in the SAP queues in the CFM, packets which are re-assembled in the ATM-capable MDAs take advantage of the application of the WRED congestion avoidance service queues associated with the ATM SAP.

  Depending of the type of MDA which supports ATM encapsulation, other MDA specific packet congestion control mechanisms operating on per SAP queues in the MDA are also applied.

Basic Configurations

A basic ATM QoS traffic descriptor profile must conform to the following:

- Each policy must have a unique policy ID.
- Default values can be modified but parameters cannot be deleted.

Create an ATM-TD-Profile QoS Policy

Configuring and applying QoS policies and profiles other than the default policy is optional.

To create an ATM QoS traffic descriptor profile, define the following:

- Assign a policy ID (policy number). The system will not dynamically assign an ID.
- Include a description. The description provides a brief overview of policy features.
- Configure traffic attributes of the ATM traffic profile.
- Determine whether egress shaping should occur.

The following displays an atm-td-profile policy configuration:

```
*A:ALA-48>config>qos>atm-td-profile# info
----------------------------------------------
            description "TEST ATM TD profile policy"
            service-category nrt-vbr
            traffic sir 4000 pir 5000
            clp-tagging
----------------------------------------------
*A:ALA-48>config>qos>atm-td-profile#
```

Applying ATM-TD-Profile Policies

Apply ATM QoS traffic descriptor profiles to the following entities:

- ATM VLL (Apipe) SAPs
- Epipe SAPs
- IES SAPs
- Ipipe SAPs
- VPRN SAPs
- VPLS SAPs

ATM VLL (Apipe) SAPs

Use the following CLI syntax to apply ATM QoS traffic descriptor profile policies to Apipe SAPs on ingress and egress.

CLI Syntax:
```
config>service>apipe>sap# atm
    egress
        traffic-desc traffic-desc-profile-id
    ingress
        traffic-desc traffic-desc-profile-id
```

Epipe SAPs

Use the following CLI syntax to apply ATM QoS traffic descriptor profile policies to Epipe SAPs on ingress and egress.

CLI Syntax:
```
config>service>epipe>sap# atm
    egress
        traffic-desc traffic-desc-profile-id
    ingress
        traffic-desc traffic-desc-profile-id
```

IES SAPs

Use the following CLI syntax to apply ATM QoS traffic descriptor profile policies to IES SAPs on ingress and egress.

CLI Syntax:
```
config>service>ies>if>sap# atm
    egress
       traffic-desc traffic-desc-profile-id
    ingress
       traffic-desc traffic-desc-profile-id
```

Ipipe SAPs

Use the following CLI syntax to apply ATM QoS traffic descriptor profile policies to Ipipe SAPs on ingress and egress.

CLI Syntax:
```
config>service>ipipe>sap# atm
    egress
       traffic-desc traffic-desc-profile-id
    ingress
       traffic-desc traffic-desc-profile-id
```

VPRN SAPs

Use the following CLI syntax to apply ATM QoS traffic descriptor profile policies to VPRN SAPs on ingress and egress.

CLI Syntax:
```
config>service>vprn>if>sap# atm
    egress
       traffic-desc traffic-desc-profile-id
    ingress
       traffic-desc traffic-desc-profile-id
```

VPLS SAPs

Use the following CLI syntax to apply ATM QoS traffic descriptor profile policies to VPLS SAPs on ingress and egress.

CLI Syntax:
```
config>service>vpls>sap# atm
    egress
       traffic-desc traffic-desc-profile-id
    ingress
       traffic-desc traffic-desc-profile-id
```

Default ATM-TD-Profile Policy Values

The default ATM QoS traffic descriptor profile is identified as `default`. The default profile cannot be edited or deleted. The following displays default profile parameters:

Table 51: ATM-TD-Profile Defaults

Field	Default
atm-td-profile *traffic-desc-profile-id*	1
description	“Default Traffic Descriptor”
service-category	ubr
traffic	no traffic
policing	no policing
clp-tagging	no clp-tagging
descriptor-type	P0_1
shaping	no shaping

The following output displays the default configuration:

```
A:ALA-48>config>qos# info detail
...
#------------------------------------------
echo "QoS Slope/Queue Policies Configuration"
#------------------------------------------
        atm-td-profile 1 create
            description "Default Traffic Descriptor"
            service-category ubr
            no traffic
            no policing
            no clp-tagging
            descriptor-type P0_1
            no shaping
        exit
        atm-td-profile 2 create
...
-------------------------------------------------
A:ALA-48>config>qos#
```

Service Management Tasks

This section discusses the following service management tasks:

- Removing a Profile from the QoS Configuration on page 670
- Copying and Overwriting Profile on page 670
- Editing QoS Policies on page 671

Removing a Profile from the QoS Configuration

The default ATM traffic descriptor profile cannot be deleted.

```
A:ALA-48>config>qos# no atm-td-profile 1
MINOR: ATM #1206 Cannot change Default Traffic Descriptor
A:ALA-48>config>qos#
```

To delete an ATM QoS traffic descriptor profile, enter the following command:

CLI Syntax: `config>qos# no atm-td-profile` *`traffic-desc-profile-id`*

Example:
```
config>qos# no atm-td-profile 2
```

Copying and Overwriting Profile

You can copy an existing profile, rename it with a new profile ID value, or overwrite an existing profile ID. The `overwrite` option must be specified or an error occurs if the destination profile ID exists.

CLI Syntax: `config>qos> copy atm-td-profile` *`src-prof dst-prof`* `[overwrite]`

Example:
```
A:ALA-48>config>qos# copy atm-td-profile 2 3
MINOR: CLI destination (3) exists use {overwrite}.
A:ALA-48>config>qos# copy atm-td-profile 2 3 overwrite
A:ALA-48>config>qos#
```

Editing QoS Policies

You can change existing policies and entries in the CLI. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors copy the policy to a work area, make the edits, and then write over the original policy.

ATM QoS Policy Command Reference

Command Hierarchies

Configuration Commands

```
config
    — qos
        — [no] atm-td-profile traffic-desc-profile-id
            — [no] clp-tagging
            — description description-string
            — no description
            — descriptor-type type
            — [no] policing
            — service-category service-category
            — [no] shaping
            — traffic [sir sir-val] [pir pir-val] [mir mir-val] [mbs mbs-val] [cdvt cdvt-val]
            — no traffic
```

Operational Commands

```
config
    — qos
        — copy atm-td-profile src-prof dst-prof [overwrite]
```

Show Commands

```
show
    — qos
        — atm-td-profile [traffic-desc-profile-id] [detail]
    — service
        — sap-using [ingress | egress] atm-td-profile td-profile-id

show
    — port [port-id] atm [detail]
    — port [port-id] atm connections [detail]
    — port [port-id] atm interface-connections [detail]
    — port [port-id] atm pvc [vpi[/vci]] [detail]
    — port [port-id] atm pvp [vpi] [detail]
    — port [port-id] atm pvt [vpi-range] [detail]
```

Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
no description

Context config>qos>>atm-td-profile

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax **copy atm-td-profile** *src-prof dst-prof* [**overwrite**]

Context config>qos

Description This command copies the source atm profile into the destination atm profile. If the destination profile was already defined, the keyword 'overwrite' must be appended for the copy to complete.

The **copy** command is a configuration level maintenance tool used to create new profiles using existing profiles. It also allows bulk modifications to an existing profile with the use of the **overwrite** keyword.

Parameters **atm-td-profile** *src-prof dst-prof* **—** Indicates that the source profile ID and the destination profile ID are atm-td-profile IDs. Specify the source ID that the copy command will copy and specify the destination ID to which the command will duplicate the profile to a new or different profile ID.

Values 1 — 1000

overwrite — Specifies to replace the existing destination profile. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, an error will occur if the destination profile ID exists.

```
A:ALA-48>config>qos# copy atm-td-profile 2 10
MINOR: CLI destination (10) exists use {overwrite}.
A:ALA-48>config>qos# copy atm-td-profile 2 10 overwrite
A:ALA-48>config>qos#
```

ATM QoS Policy Commands

atm-td-profile

Syntax [**no**] **atm-td-profile** *traffic-desc-profile-id*

Context config>qos

Description This command is used to configure an ATM traffic descriptor profile.

Traffic descriptor profiles are used to:

1. Define traffic management capabilities for ATM PVCCs.
2. Calculate the total bandwidth consumed on a given port by all ATM PVCC(s). The BW taken by a PVCC is equal to:

   a. PIR for CBR PVCCs

   b. SIR for rt-vbr and nrt-vbr PVCCs

   c. MIR for UBR PVCC
3. Define ATM-level SAR scheduling

The default traffic descriptor is pre-configured and non-modifiable. It cannot be deleted. All other traffic descriptor profiles must be explicitly created before use. The create keyword must follow each new profile configuration.

Any changes made to the existing profile, using any of the sub-commands are applied immediately to all objects where this profile is applied (a small traffic interruption in data traffic will occur during the data plane reprogramming with the newly modified profile).

When many changes are required on a profile, it is recommended that the profile be copied to a work area profile ID. That work-in progress profile can be modified until complete and then written over the original profile-id. Use the config qos copy command to maintain profiles in this manner.

The weight assigned to each non-shaped PVCC in the Deficit Round Robin Scheduler depends on the service category and traffic rates (see traffic command for more details).

The **no** form of the command deletes a given traffic profile. Note that the profile to be deleted must not be associated with any object (for example a SAP). If this condition is not met, the command will return an error.

Default 1 — Default Traffic Descriptor (UBR, no traffic, no shaping)

Parameters *traffic-desc-profile-id* — Index identifier for a traffic descriptor profile

Values 1 — 1000

clp-tagging

Syntax [**no**] **clp-tagging**

Context config>qos>atm-td-profile

Description This command controls the setting of the CLP bit in the ATM cell header for egress traffic on an IES or VPRN SAP.

When enabled, traffic queued on expedited queues has the CLP bit set to zero, while traffic on non-expedited queues has the CLP bit set to one.

The **no** form of the command sets the CLP bit set to zero.

Default no clp-tagging

descriptor-type

Syntax **descriptor-type** {*type*}

Context config>qos>atm-td-profile

Description This command is used to specify the type of the traffic descriptor profile as per ATM Forum Traffic Management Specification Version 4.1.

Parameters **Values** P0_1, P0_1andS0_Tag, P0_1andS0, P0_1andS0_1

The descriptor type defines interpretation of traffic parameters that are specified for this profile. The following table details these rules:

Descriptor Type	Rates Interpretation	Applicable Service Categories
P0_1	PIR applies to CLP=0 and CLP=1 cell flows	CBR, UBR, UBR with MIR
P0_1andS0_1	PIR applies to CLP=0 and CLP=1 cell flows SCR applies to CLP=0 and CLP=1 cell flows	rt-VBR and nrt-VBR
P0_1andS0	PIR applies to CLP=0 and CLP=1 cell flows SCR applies to CLP=0 cell flow	rt-VBR and nrt-VBR

Setting descriptor type to a value not compatible with the service category (as defined in the above table) is an error.

Default The following table defines default values of descriptor type based on a service category:

Service Category	Default Descriptor Type
CBR	P0_1
UBR	P0_1
UBR with MIR	P0_1
rt-VBR or nrt-VBR	P0_1andS0_1

policing

Syntax [**no**] **policing**

Context config>qos>atm-td-profile

Description This command determines whether ingress traffic is policed. Policing is valid for CBR, RT-VBR and NRT-VBR. This is cell-based policing.

Default disabled

service-category

Syntax **service-category** *service-category*

Description config>qos>atm-td-profile

Description This command is used to configure an ATM service category attribute of an ATM traffic descriptor profile per ATM Forum Traffic Management Specification Version 4.1.

Parameters The router supports the following ATM service categories on ATM-capable MDAs:

Service Category	Description
CBR	Constant Bit Rate
rt-VBR	real time Variable Bit Rate
nrt-VBR	non-real time Variable Bit Rate
UBR	Unspecified Bit Rate without Minimum Desired Cell Rate (defined by specifying service category to be ubr, and MIR of 0)
UBR (with MIR)	Unspecified Bit Rate with non-zero Minimum Desired Cell Rate (defined by specifying service category to be ubr, and MIR > 0)

Changing the service category of a profile will reset all traffic attributes to their defaults (see the traffic command) and will cause reprogramming of the data path (with a small impact on user traffic) and a reset of VC statistics for all VCs using this traffic descriptor profile.

Default ubr

shaping

Syntax [**no**] **shaping**

Context config>qos>atm-td-profile

Description This command enables cell level shaping when the ATM traffic descriptor profile is applied to an ATM SAP queue. Shaping is only applied in the egress queue of the ATM SAP. Shaping cannot be enabled on an ATM SAP with the UBR service category.

The **no** form of this command disables shaping.

Default The default is determined by the service category. The following default applies for shaping depending upon a given service category:

Applicable Service Category	Default Shaping Value	Comments
UBR	disabled	Shaping cannot be enabled
CBR	enabled	Shaping cannot be disabled when the profile is applied to ATM SAP on ATM MDA
rt-VBR	enabled	Shaping cannot be disabled when applied to ATM SAP on ATM MDA
nrt-VBR	enabled	

traffic

Syntax **traffic** [**sir** *sir-val* [**pir** *pir-val*] [**mir** *mir-val*] [**mbs** *mbs-val*] [**cdvt** *cdvt-val*]
no traffic

Context config>qos>atm-td-profile

Description This command is used to configure traffic attributes of an ATM traffic profile as per ATM Forum Traffic Management Specification Version 4.1.

The traffic parameters of a traffic descriptor that are configurable depends on the service category of this traffic descriptor profile (see the service-category command).

The following table defines which traffic descriptor parameters are applicable for what service category and what are configuration rules between the parameters. **Y** indicates the parameter can be configured for a given service category and will be defaulted if not provided, an **N/A** indicates the parameter cannot be configured for a given service category (an error will be returned). If an applicable parameter is not specified, the current value will be preserved.

Service Category	SIR	PIR	MBS	MIR	CDVT
CBR	N/A	Y	N/A	N/A	Y
rt-VBR	Y	Y (must be >=SIR)	Y	N/A	Y
Nrt-VBR	Y	Y (must be >=SIR)	Y	N/A	Y
UBR	N/A	Y	N/A	N/A	N/A
UBR with MIR	N/A	Y (must be >= MIR	N/A	Y (non-zero MIR specified)	N/A

Configuring PIR for traffic descriptor profiles for UBR and UBR with MIR service categories has no impact on a traffic contract when a PVCC using that profile resides on an m4-atmoc12/3-sfp MDA. On this MDA SAR ignores PIR (de-facto treating each UBR as it would have a PIR of max. line rate). The default pir value for UBR and UBR with MIR reflects this behavior.

When a traffic descriptor profile is used to define egress scheduling, the following describes how traffic rates are used to derive scheduling weight:

1. UBR PVCCs (i.e., MIR = 0) are assigned weight value of 1
2. UBR with MIR PVCCs are assigned weight value in the inclusive range from 1 to 255 based on the MIR rate.
3. rt-VBR and nrt-VBR PVCCs are assigned weight value in the inclusive range from 1 to 255 based on the SCR rate
4. CBR PVCCs are assigned weight value in the inclusive range from 1 to 255 based on the PIR rate

The scheduling weight is derived from the traffic rate based on the following formula:

If traffic rate <= 32 Kbps, then weight = 1

If 32 Kbps < traffic rate < 8160 Kbps, then weight = floor (traffic rate / 32)

If traffic rate >= 8160 Kbps, then weight = 255

The configuration of weight unit (32 Kbps) is left for future releases.

Since the SAR operates in cells/second with 1 cell granularity, PIR and SCR values programmed need to be converted to cells per second. When converting values to be used for scheduler, the result is rounded up to the next cell when required by conversion.

When any of SIR, PIR, or MIR is greater than the physical maximum port/channel capacity for a given PVCC, then the maximum physical port/channel capacity is used in BW accumulation and when configuring the H/W for that PVCC.

Hardware- enforceable mbs is in the inclusive range from 3 to 256 000 cells. Any value outside of that range will be accepted and rounded up/down to the minimum/maximum enforceable value.

The **no** form of the command restores traffic parameters to their defaults for a given service category.

By default ATM traffic parameters are, in kbps:

Service Category	Traffic Parameter Defaults
CBR:	
PIR	0
rt-VBR and nrt-VBR	
PIR	0
SCR	0
MBS	32
UBR (note by default UBR is without MIR)	
PIR	0
MIR	0

Parameters **sir** *value* **—** Sustained Information Rate (including cell overhead) in kilobits per second.

Values 0 — 4294967295

pir *value* **—** Peak Information Rate (including cell overhead) in kilobits per second.

Values 0 — 4294967295

mir *value* **—** Minimum Desired Information Rate (including cell overhead) in kilobits per second.

Values 0 — 4294967295

mbs *value* **—** Maximum Burst Size in cells

Values 0 — 4294967295

cdvt *cdvt-val* **—** "The Cell Delay Variation Tolerance (CDVT), in microseconds.

Default Depending upon a given service category:

CBR/RT-VBR/NRT-VBR 250

Values 0 — 4294967295

Show Commands

atm-td-profile

Syntax **atm-td-profile** [*traffic-desc-profile-id*] [**detail**]

Context show>qos

Description This command displays ATM traffic descriptor profile information.

Parameters *traffic-desc-profile-id* — Displays the ATM traffic descriptor profile.

Values 1 — 1000

detail — Displays detailed policy information including policy associations.

Output **ATM TD Profile Output —** The following table describes ATM traffic descriptor profile show command output.

Label	Description
Maximum Supported Profiles	Displays the maximum number of ATM traffic descriptor profiles that can be configured on this system.
Currently Configured Profiles	Displays the number of currently configured ATM traffic descriptor profiles on this system.
TDP-Id	The ID that uniquely identifies the traffic descriptor policy.
Description	A text string that helps identify the policy's context in the configuration file.
Service Category	Displays the ATM service category.
SCR	Displays the sustained cell rate in Kbps.
PIR	Displays the peak cell rate in Kbps.
MIR	Displays the Minimum Desired Cell Rate in Kbps.
MBS	Displays the maximum burst size in cells.
Shaping	Displays whether shaping is enabled or disabled for the traffic descriptor profile.
Entities using TDP-ID	Displays the number of entities using the ATM traffic descriptor.
-	Indicates that the parameter is not applicable for the configured service category.

```
A:ALA-48>config>qos>atm-td-profile# show qos atm-td-profile
===============================================================================
Traffic Descriptor Profiles
===============================================================================
Maximum Supported Profiles     : 1000
Currently Configured Profiles  : 3
-------------------------------------------------------------------------------
TDP-id Description
       Service Category SCR         PIR         MIR        MBS
-------------------------------------------------------------------------------
1      Default Traffic Descriptor
       UBR              -           0           0           -
-------------------------------------------------------------------------------
2      Default Traffic Descriptor
       NRT_VBR          4000        5000        -           32
-------------------------------------------------------------------------------
10     Default Traffic Descriptor
       NRT_VBR          4000        5000        -           32
===============================================================================
A:ALA-48>config>qos>atm-td-profile#

A:ALA-48>config>qos>atm-td-profile# show qos atm-td-profile 10 detail
===============================================================================
Traffic Descriptor Profile (10)
===============================================================================
-------------------------------------------------------------------------------
TDP-id Description
       Service Category SCR         PIR         MIR        MBS
-------------------------------------------------------------------------------
10     Default Traffic Descriptor
       NRT_VBR          4000        5000        -           32
-------------------------------------------------------------------------------
TDP details
-------------------------------------------------------------------------------
Shaping       : disabled
-------------------------------------------------------------------------------
Entities using TDP-10
-------------------------------------------------------------------------------
===============================================================================
A:ALA-48>config>qos>atm-td-profile#
```

sap-using

Syntax **sap-using** [**ingress** | **egress**] **atm-td-profile** *td-profile-id*

Context show>service

Description Displays **atm-td-profile** SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **ingress —** Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

qos-policy *qos-policy-id* **—** The ingress or egress QoS Policy ID for which to display matching SAPs.

Values 1 — 65535

filter *filter-id* **—** The ingress or egress Filer Policy ID for which to display matching SAPs.

Values 1 — 65535

sap-id — Specifies the physical port identifier portion of the SAP definition.

Values: *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id*][:*vpi/vci*|*vpi*| *vpi1.vpi2*]	
frame	[*port-id* | *aps-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
cem	*slot/mda/port.channel*	
ima-grp	[*bundle-id*[:vpi/vci|vpi|*vpi1.vpi2*]	
port-id	*slot/mda/port*[.*channel*]	
bundle-id	bundle-*type-slot/mda.bundle-num*	
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 256
bpgrp-id	bpgrp-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima, ppp
	bpgrp-num	1 — 1280
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 16
lag-id	lag-*id*	
	lag	keyword
	id	1 — 64
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI: 0 — 4095	
	UNI: 0 — 255	

vci	1, 2, 5 — 65535
dlci	16 — 1022

interface — Specifies matching SAPs with the specified IP interface.

ip-addr — The IP address of the interface for which to display matching SAPs.

Values 1.0.0.0 — 223.255.255.255

ip-int-name — The IP interface name for which to display matching SAPs.

td-profile-id — Profile ID that identifies a specific profile to display.

Output **Show Service SAP —** The following table describes show service SAP output fields:

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The service identifier.
SapMTU	The SAP MTU value.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
E.QoS	The SAP egress QoS policy number specified on the egress SAP.
E.Mac/IP	The MAC or IP filter policy ID applied to the egress SAP.
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The desired state of the SAP.
Opr	The actual state of the SAP.

Sample Output

```
A:ALA-48>config>service>ies# show service sap-using sap 1/3/2:244/1
===============================================================================
Service Access Points Using Port 1/3/2:15990785
===============================================================================
PortId                SvcId      I.QoS I.Fltr E.QoS E.Fltr A.Pol  Adm  Opr
-------------------------------------------------------------------------------
1/3/2:244/1           89         1     none   1     none   none   Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
A:ALA-48>config>service>ies#
```

port

Syntax **port** [*port-id*] **atm**
port [*port-id*] **atm connections**
port [*port-id*] **atm interface-connections**
port [*port-id*] **atm pvc**
port [*port-id*] **atm pvp**
port [*port-id*] **atm pvt**

Context show

Description This command displays port or channel information.

Parameters *port-id* — Specifies the physical port ID in the form *slot/mda/port*.

Syntax *slot*[/*mda*[/*port*[.*sonet-sdh-index*]]]

Slot Values 7710 SR-c4: 1

MDA Values 1, 3, 5, 7, 9, 11 (SR-c12), 1, 3 (SR-c4)

CMA Values 1 — 12 (SR-c12), 1 — 4 (SR-c4)

Port Values 1 — 60 (depending on the MDA)

Channelized Port Values (for channelized MDAs):

M1-CHOC12-SFP:	slot/mda/port. [1..4] . [1..3] . [1..28] . [..24] For example, 1/2/1.1.1.28.24
M12-DS3:	slot/mda/port. [1..28] . [..24] For example, 1/1/1.1.1

connections — Display ATM connection information

interface-connection — Display ATM interface connection information

pvc — Displays ATM port PVC information

pvp — Displays ATM port PVP information

pvt — Displays ATM port PVT information

vpi — Specifies the ATM network virtual path identifier (VPI) for this PVC.

vci — Specifies the ATM network virtual channel identifier (VCI) for this PVC.

detail — Provides detailed information.

Output **Port ATM PVC Detail Output —** The following table describes port ATM PVC detail output fields.

Table 52: Show Port ATM PVC VPI/VCI Detail Output Fields

Label	Description
Port Id	The port ID configured or displayed in the *slot/mda/port* format.
VPI/VCI	Displays the VPI/VCI values.
Admin State	Displays the administrative state of the interface connection.
Oper State	Indicates the status of the ATM interface.

Table 52: Show Port ATM PVC VPI/VCI Detail Output Fields (Continued)

Label	Description
OAM State	Indicates the OAM operational status of ATM connections. ETE indicates end-to-end connection. AIS denotes alarm indication signal. RDI denotes for remote defect indication. AIS-LOC indicates the alarm was due to loss of continuity of periodic loopbacks.
Encap Type	Indicates the encapsulation type.
Owner	Identifies the system entity that owns a specific ATM connection.
AAL Type	Displays ATM Adaptation Layer 5 (AAL5) information.
Endpoint Type	Displays the endpoint type.
Cast Type	Indicates the connection topology type.
Type	Indicates the connection type.
Ing. Td Idx	Specifies the ATM traffic descriptor profile that applies to the receive direction of the interface connection.
Egr. Td Idx	Specifies the ATM traffic descriptor profile that applies to the transmit direction of the interface connection.
Last Changed	Indicates the date and time when the interface connection entered its current operational state.
Octets	Displays the number of input and output octets. HEC discarded cells are not included in the input octet numbers.
Cells	Displays the number of input and output cells. HEC discarded cells are not included in the input cell numbers.
Packets	Displays the number of input and output packets. Packets discarded due to HEC or oversize discards are not counted. CRC errored are also in the packet counts and display on the VC level statistics but not on the port level.
Dropped Packets	Displays the number of packets dropped by the ATM SAR device.
CRC-32 Errors	Displays the number of valid AAL-5 SDUs and AAL-5 SDUs with CRC-32 errors received by the AAL-5 VCC.
Reassembly Time-outs	Displays the number of reassembly timeout occurrences.
Over Sized SDUs	Displays the total number of oversized SDU discards.
AIS	Displays the number of AIS cells transmitted and received on this connection for both end to end and segment.
RDI	Displays the number of RDI cells transmitted and received on this connection for both end to end and segment.

Table 52: Show Port ATM PVC VPI/VCI Detail Output Fields (Continued)

Label	Description
Loopback	Displays the number of loopback requests and responses transmitted and received on this connection for both end to end and segment.
CRC-10 Errors	Displays the number of cells discarded on this VPL with CRC 10 errors.
Other	Displays the number of OAM cells that are received but not identified.

Sample Output

```
A:ALA-1# show port 1/1/2 atm pvc 0/500 detail
===============================================================================
ATM Endpoint
===============================================================================
Port Id            : 1/1/2               VPI/VCI            : 0/500
Admin State        : up                  Oper state         : down
OAM State          : ETE-AIS             Encap Type         : llc
Owner              : SAP                 AAL Type           : AAL-5
Endpoint Type      : PVC                 Cast Type          : P2P
Ing. Td Idx        : 5                   Egr. Td Idx        : 3
Last Changed       : 02/14/2007 14:15:12
===============================================================================
ATM Statistics
===============================================================================
                                                    Input                 Output
-------------------------------------------------------------------------------
Octets                                                  0                      0
Cells                                                   0                      0
===============================================================================
AAL-5 Packet Statistics
===============================================================================
                                                    Input                 Output
-------------------------------------------------------------------------------
Packets                                                 0                      0
Dropped Packets                                         0                      0
CRC-32 Errors                                           0
Reassembly Timeouts                                     0
Over Sized SDUs                                         0
===============================================================================
ATM OAM Statistics
===============================================================================
                                                    Input                 Output
-------------------------------------------------------------------------------
AIS                                                     0                      0
RDI                                                     0                      0
Loopback                                                0                      0
CRC-10 Errors                                           0
Other                                                   0
===============================================================================
A:ALA-1#
```

QoS in MC-MLPPP

In This Section

This section provides information to configure MC-MLPPP using the command line interface.

Topics in this section include:

- Overview on page 692
- Basic Configurations on page 697
- Configuring MC-MLPPP on page 698

Overview

If the user enables the multiclass option under an MLPPP bundle, the MDA egress data path provides a queue for each of the 4 classes of MLPPP. The user configures the required number of MLPPP classes to use on a bundle. The forwarding class of the packet, as determined by the ingress QoS classification, is used to determine the MLPPP class for the packet and hence which of the four egress MDA queues to store the packet. The mapping of forwarding class to MLPPP class is a function of the user configurable number of MLPPP classes. The default mapping for a 4-class, 3-class, and 2-class MLPPP bundle is shown in Table 53.

Table 53: Default Packet Forwarding Class to MLPPP Class Mapping

FC ID	FC Name	Scheduling Priority (Default)	MLPPP Class 4-class bundle	MLPPP Class 3-class bundle	MLPPP Class 2-class bundle
7	NC	Expedited	0	0	0
6	H1	Expedited	0	0	0
5	EF	Expedited	1	1	1
4	H2	Expedited	1	1	1
3	L1	Non-Expedited	2	2	1
2	AF	Non-Expedited	2	2	1
1	L2	Non-Expedited	3	2	1
0	BE	Non-Expedited	3	2	1

Table 54 shows a different mapping enabled when the user applies one of three pre-defined egress QoS profiles in the 4-class bundle configuration only.

Table 54: Packet Forwarding Class to MLPPP Class Mapping

FC ID	FC Name	Scheduling Priority (Default)	MLPPP Class (MLPPP Egress QoS profile 1, 2, and 3)
7	NC	Expedited	0
6	H1	Expedited	0
5	EF	Expedited	1
4	H2	Expedited	2
3	L1	Non-Expedited	2
2	AF	Non-Expedited	2
1	L2	Non-Expedited	2
0	BE	Non-Expedited	3

The MLPPP class queue parameters and its scheduling parameters are also configured by applying one of the three pre-defined egress QoS profiles to an MLPPP bundle.

Table 55 and Figure 30 provide the details of the class queue threshold parameters. Packets marked with a high drop precedence, such as out-of-profile, by the service or network ingress QoS policy will be discarded when any class queue reaches the OOP threshold. Packet with a low drop precedence marking, such as in-profile, will be discarded when any class queue reaches the max threshold.

Table 55: MLPPP Class Queue Threshold Parameters

	Class 0		Class 1		Class 2		Class 3	
Queue Threshold (in ms @ Available bundle rate)	Max	Oop	Max	Oop	Max	Oop	Max	Oop
2-Class Bundle Default Egress QoS Profile	250	125	750	375	N/A	N/A	N/A	N/A
3-Class Bundle Default Egress QoS Profile	50	25	200	100	750	375	N/A	N/A
4-Class Bundle Default Egress QoS Profile	10	5	50	25	150	75	750	375
4-Class Bundle Egress QoS Profile 1	25	12	5	3	200	100	1000	500
4-Class Bundle Egress QoS Profile 2	25	12	5	3	200	100	1000	500
4-Class Bundle Egress QoS Profile 3	25	12	5	3	200	100	1000	500

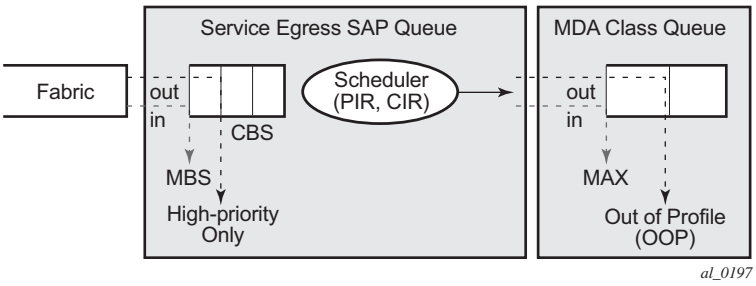

Figure 30: MLPPP Class Queue Thresholds for In-Profile and Out-of-Profile Packets

Table 56 and Table 31 provide the details of the class queue scheduling parameters.

Table 56: MLPPP Class Queue Scheduling Parameters

		WRR Parameters		
4-class MLPPP Egress QoS Profile	MIR	W1	W2	W3
Profile 1	85%	<1%	66%	33%
Profile 2	90%	<1%	89%	10%
Profile 3	85%	<1%	87%	12%

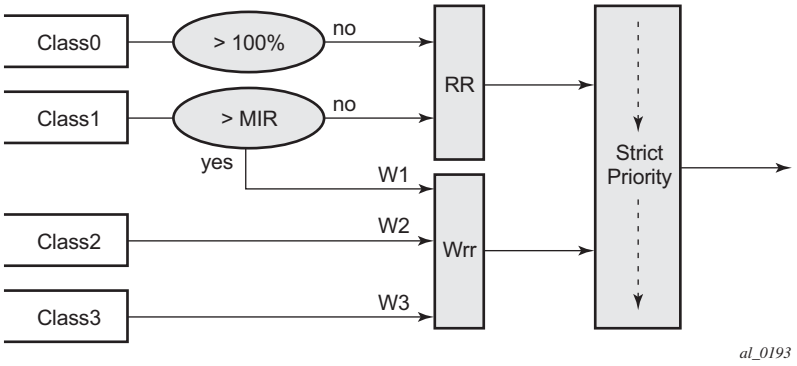

Figure 31: MLPPP Class Queue Scheduling Scheme

Note that all queue threshold and queue scheduling parameters are adjusted to the available bundle rate. If a member link goes down or a new member link is added to the bundle, the scheduling parameters MIR, W1, W2, W3, as well as the per class queue thresholds OOP and max are automatically adjusted to maintain the same values.

Class 0 queue is serviced at MLPPP at available bundle rate. Class 1 queue is guaranteed a minimum service rate but is allowed to share additional bandwidth with class 2 and 3 queues based on the configuration of WRR weight W1.

Class queues 2 and 3 can be given bandwidth guarantee by limiting MIR of class 1 queue to less than 100% and by setting the WRR weights W1, W2 and W3 to achieve the desired bandwidth distribution among all three class queues.

Note that there is one queue per bundle member link to carry link control packets, such as LCP: PPP, and which are serviced with strict priority over the 4 class queues (not shown).

In the default 2-class, 3-class, and 4-class egress QoS profile, the class queues are service with strict priority in ascending order of class number.

Ingress MLPPP Class Reassembly

For a MLPPP bundle with the multi-class option enabled, there is a default profile for setting the re-assembly timer value for each class. When the pre-defined MLPPP ingress QoS profile 1 is applied to a 4-class bundle, the values of the timers are modified as shown in Table 57.

Table 57: MLPPP Ingress QoS Profile: Reassembly Timers (msec)

	Class 0	Class 1	Class 2	Class 4
MLPPP ingress QoS default profile (2-Class bundle)	25ms	25ms	NA	NA
MLPPP ingress QoS default profile (3-Class bundle)	25ms	25ms	25ms	NA
MLPPP ingress QoS default profile (4-Class bundle)	25ms	25ms	100ms	1000ms
MLPPP ingress QoS profile 1 (4-class bundle)	10	10	100	1000

Configuring MC-MLPPP QoS Parameters

A 4-class MLPPP bundle can be configured. This feature cannot be used with MC-MLPPP bundles with fewer than 4 classes.

The following describe the parameters and the configuration processes and rules:

1. The user creates an ingress QoS profile, mlppp-profile-ingress, to configure the desired value of the ingress per-class re-assembly timer. Ingress QoS profile #1 is reserved to the pre-defined profile with parameter values in Table 57. The user is allowed to edit this profile and change parameter values. However, the default value of a parameter when a user creates a profile with a profile-id higher than 1, or performs the no option on the parameter, will always be the one in Table 57 for the ingress QoS Profile #1 regardless what parameter value the edited Profile #1 has at that point in time.
2. The user creates an egress QoS profile, mlppp-profile-egress, to configure the desired values for the per-class queue and queue scheduling parameters. The user is also able to configure the mapping of the system forwarding classes to the MLPPP classes. Egress QoS profiles #1, 2, and 3, are reserved to the pre-defined profiles with parameter values shown in Table 54, Table 55, or Table 56. The user is allowed to edit these profiles and change parameter values. However, the default value of a parameter when a user creates a profile with a profile-id

higher than 3, or when the user performs the no option on the parameter, will be the one shown in Table 54, Table 55, or Table 56 for the egress QoS Profile #1. This is regardless of the parameter value the edited profiles have at that point in time.

3. A maximum of 128 ingress QoS profiles and 128 egress QoS profiles can be created on the system.
4. The values of the ingress per-class re-assembly timer are configured in the ingress QoS profile.
5. The mapping of the system forwarding classes to the MLPPP Classes are configured in the egress QoS profile. There is a many-to-one relationship between the system FC and an MLPPP class. See Table 54 for the mapping when one of the three pre-defined 4-class egress QoS profiles is selected.
6. The maximum size for each MLPPP class queue in units of msec at the available bundle rate is configured in the egress QoS profile. This is referred to as max in Figure 30 and as max-queue-size in CLI. The out-of-profile threshold for an MLPPP class queue, referred to as oop in Figure 30, is not directly configurable and is set to 50% of the maximum queue size rounded up to the nearest higher integer value.
7. The MLPPP class queue scheduling parameters is configured in the egress QoS profile. The minimum information rate, referred to as MIR in Figure 31 and mir in CLI, applies to Class 1 queue only. The MIR parameter value is entered as a percentage of the available bundle rate. The WRR weight, referred to as W1, W2, and W3 in Figure 31 and weight in CLI, applies to class 1, class 2, and class 3 queues. Note that W1 in Figure 31 is not configurable and is internally set to a value of 1 such that Class 1 queue shares 1% of the available bundle rate when the sum of W1, W2, and W3 equals 100. W2 and W3 weights are integer values and are user configurable such that Class 2 queue shares and Class 3 queue shares of the available bundle rate.
8. The user applies the ingress and egress QoS profiles to a 4-class MLPPP bundle for the configured QoS parameter values to take effect on the bundle.
9. The following operations require the bundles associated with a QoS profile to be shutdown to take effect.
   - → A change of the numbered ingress or egress QoS profile associated with a bundle.
   - → A change of the bundle associated ingress or egress QoS profile from default profile to a numbered profile and vice-versa.
10. The following operations can be performed without shutting down the associated bundles:
   - → Changes to any parameters in the ingress and egress QoS profiles.

Basic Configurations

Configure an egress and ingress MLPPP profile.

The following displays the profile configuration examples:

```
A:ALA-12>config>qos# info
#------------------------------------------
         mlppp-profile-ingress 2 [create]
             description my-4-class-bundle-ingress-profile
             class class 0
                      reassembly-timeout 10
                  class class 1
                      reassembly-timeout 100
                  class class 2
                      reassembly-timeout 500
                  class class 3
                      reassembly-timeout 1000
         mlppp-profile-egress 4 [create]
            description my-4-class-bundle-egress-profile
             fc be mlppp-class 3
             fc l2 mlppp-class 2
             fc af mlppp-class 2
             fc l1 mlppp-class 2
             fc h2 mlppp-class 2
             fc ef mlppp-class 1
             fc h1 mlppp-class 1
             fc nc mlppp-class 0
                   class 0
                  max-queue-size  10
             class 1
                  max-queue-size  50
                  mir 25
                    class 2
                  max-queue-size  100
                  weight 25
                    class 3
                  max-queue-size  1000
                  weight 75

#------------------------------------------
A:ALA-12>config>qos#


A:ALA-12>config>port>ml-bundle# info
#------------------------------------------
              multilink-bundle bundle-ppp-6/1.1
                multilink-bundle
                   fragment-threshold 384
                   mlppp
                  multiclass 4
                           ingress
                              qos-profile 2
                     exit
                           egress
                              qos-profile 4
                           exit
                      exit
```

```
            exit
            member 1/1/1.1.1.1
            member 1/1/1.1.2.1
            member 1/1/1.1.3.1
            member 1/1/1.1.4.1
            minimum-links 2
```

Configuring MC-MLPPP

Use the following CLI syntax to create an MC-MLPPP:

```
CLI Syntax:config>qos# mlppp-profile-ingress 2 [create]
            description dmy-4-class-bundle-ingress-profile
            class class 0
               reassembly-timeout 10
            class class 1
               reassembly-timeout 100
            class class 2
               reassembly-timeout 500
            class class 3
               reassembly-timeout 1000

CLI Syntax:config>qos# mlppp-profile-egress 4 [create]
            description dmy-4-class-bundle-egress-profile
```

QoS in MLFR and FRF.12 Fragmentation

The following sections describe MLFR and FRF.12 feature descriptions and implementation.

QoS in MLFR

The MLFR feature introduces the following new QoS requirements on the MDA:

- Four MDA queues are provided per MLFR bundle to store the fragments of the FR SAP packets. One queue per FR scheduling class. Fragments of all FR SAPs of a given scheduling class will be queued in the same queue.
- The fragments of an FR SAP packet must be queued in the same MDA queue regardless of which forwarding class queue they use in the IOM.

The FR class queue parameters and its scheduling parameters are configured by applying an Egress QoS profile to an MLFR bundle.

Table 58 and Figure 32 provide the details of the class queue threshold parameters. Packets that are marked with high drop precedence, for example, out-of-profile, by the service or network ingress QoS policy will be discarded when any class queue reaches the OOP threshold. Packets with a low drop precedence marking, for example, in profile, will be discarded when any class queue reaches the max threshold. Only the max threshold is user configurable and is referred to as max-queue-size in the CLI. The OOP threshold is always set to 50% of the max threshold.

Table 58: Default FR Class Queue Threshold Parameters

	Class 0		Class 1		Class 2		Class 3	
Queue threshold (in ms@available bundle rate	Max	Oop	Max	Oop	Max	Oop	Max	Oop
	10	5	50	25	150	75	750	375

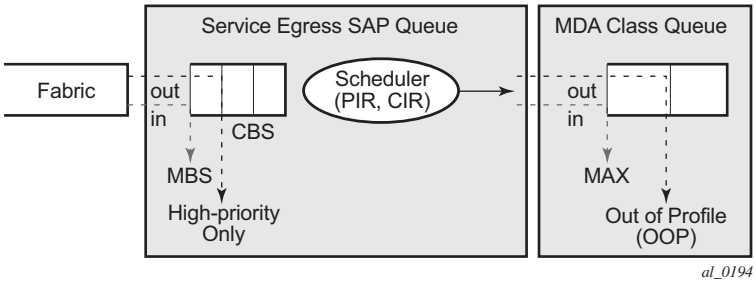

Figure 32: FR Class Queue Thresholds for In-Profile and Out-of-Profile Packets

Table 59 and Figure 33 provide the details for the class queue scheduling parameters for an MLFR bundle.

Table 59: Default FR Class Queue Scheduling Parameters

	WRR Parameters		
MIR	W1	W2	W3
90%	<1%	89%	10%

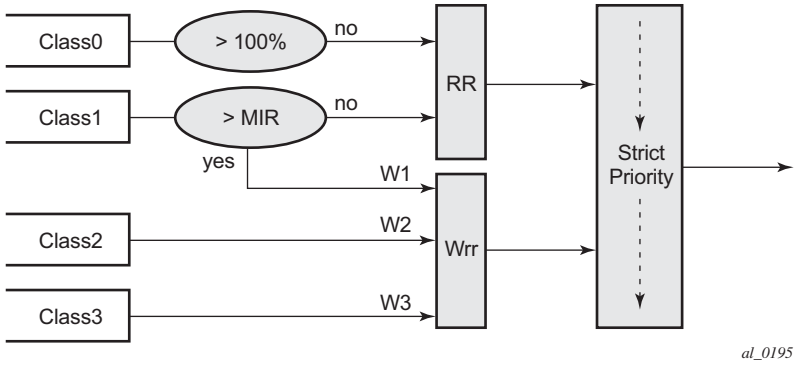

Figure 33: FR Class Queue Scheduling for an MLFR Bundle

The minimum information rate, referred to as MIR in Figure 33 and MIR in CLI, applies to Class 1 queues only. The MIR parameter value is entered as a percentage of the available bundle rate. The WRR weight, referred to as W1, W2, and W3 in Table 59 and weight in CLI, applies to class 1, class 2, and class 3 queues. Note that W1 is not configurable and is internally set to a value of 1 such that Class 1 queue shares 1% of the available bundle rate when the sum of W1, W2, and W3 equals 100. W2 and W3 weights are integer values and are user configurable such that Class 2

queue shares W2/(W1+W2+W3) and Class 3 queue shares W3/(W1+W2+W3) of the available bundle rate.

Note that all queue threshold and queue scheduling parameters are adjusted to the available bundle rate. If a member link goes down or a new member link is added to the bundle, the scheduling parameters MIR, W1, W2, W3, as well as the per class queue thresholds OOP and max are automatically adjusted to maintain the same values.

In addition, operator user can configure the value of the FR scheduling class ingress re-assembly timeout for an MLFR bundle. The default values of the timers are shown in Table 60.

Table 60: Default FR Ingress QoS Profile: Reassembly Timers (msec)

Class 0	Class 1	Class 2	Class 3
10	10	100	100

The following operations require the bundles or links associated with a QoS profile to be shutdown to take effect.

- A change of the numbered ingress or egress QoS profile.

The following operations can be performed without shutting down the associated bundles or links:

- Changes of parameters in the currently assigned ingress and egress QoS profiles.

QoS in FRF.12 End-to-End Fragmentation

When end-to-end fragmentation is enabled on an FR SAP, the queuing and scheduling on the MDA is reusing the existing FR SAP high and low priority queues are show in Figure 34:

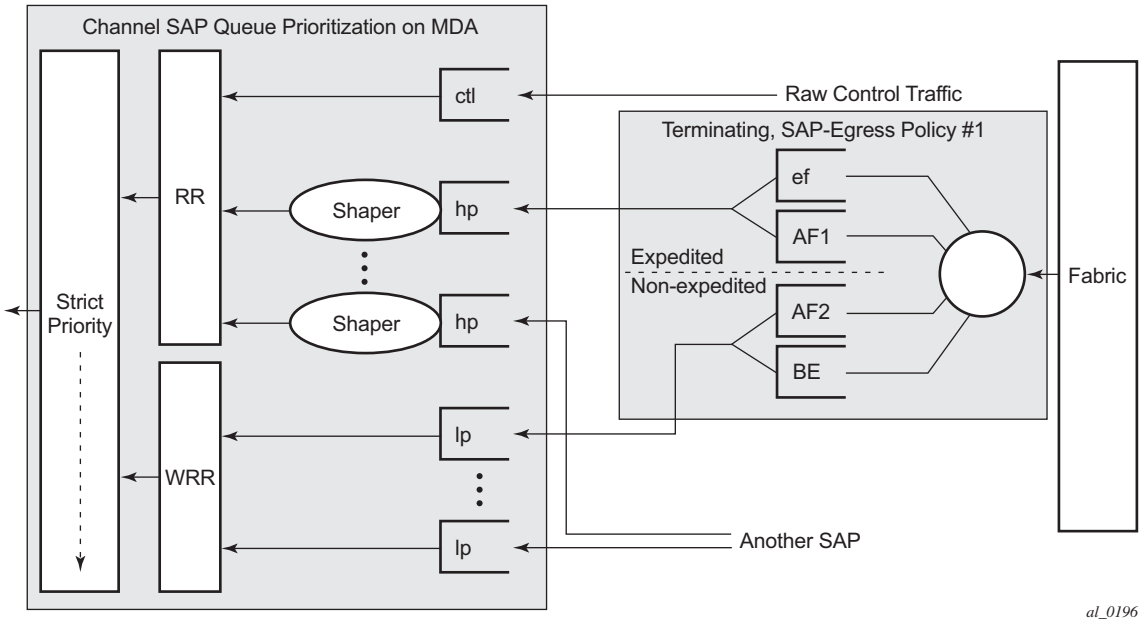

Figure 34: DLC Egress Channel Queue Scheduling

Some minor modifications are introduced to accommodate end-to-end FRF.12 fragments. The user configured FR scheduling class for this SAP will dictate which of the FR SAP MDA queues should be used to queue the fragments. Classes 0 and 1 map all fragments of the FR SAP to the Hi-Priority SAP queue. Classes 2 and 3 map all fragments of the FR SAP to the Low-Priority SAP queue.

The scheduling parameters of these queues have to be modified from existing ones as follows:

- Hi priority (HP) SAP queue: SAP_HP_Q_PIR= Sum{all SAP_FC_Q_PIR}
- Low priority (LP) SAP queue: SAP_LP_Q_WRR_Weight = Sum{all SAP_FC_Q_CIR}

MLPPP Command Reference

Command Hierarchies

Configuration Commands

```
config
    — qos
        — mlppp-profile-egress profile-id [create]
        — no mlppp-profile-egress profile-id
            — class class-id
                — max-queue-size queue-size
                — no max-queue-size
                — mir mir-value
                — no mir
                — weight weight-value
                — no weight
            — description description-string
            — no description
            — fc fc-class mlppp-class class-id
            — no fc fc-class
        — mlppp-profile-ingress profile-id [create]
        — no mlppp-profile-ingress profile-id
            — class class-id
                — reassembly-timeout timeout-value
                — no reassembly-timeout
            — description description-string
            — no description
```

Configuration Commands

MC-MLPPP Commands

mlppp-profile-ingress

Syntax [**no**] **mlppp-profile-ingress** *profile-id* [**create**]

Context config>qos

Description This command creates a profile for the user to configure the ingress QoS parameters of a multi-class MLPPP bundle.

A maximum of 128 ingress QoS profiles can be created on the system.

The **no** form of this command deletes the profile.

Parameters *profile-id* — Specifies a specific multi-class ingress profile.

Values 1 — 65535

class

Syntax **class** *class-id*

Context config>qos>mlppp-profile-ingress
config>qos>mlppp-profile-egress

Description This command provides the MLPPP class context for the user to configure the ingress or egress MLPPP bundle QoS parameters for this profile.

Parameters *class-id* — Specifies a class for this policy.

Values 0— 3

reassembly-timeout

Syntax **reassembly-timeout** *timeout-value*
no reassembly-timeout

Context config>qos>mlppp-profile-ingress>class

Description This command configures the value of the MLPPP bundle ingress per class reassembly timer for this profile.

Parameters *timeout-value* — Specifies a a reassembly timeout for this policy.

Values 1 — 1000 in milliseconds

description

Syntax **description** *description-string*
no description

Context config>qos>mlppp-profile-ingress
config>qos>mlppp-profile-egress

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

mlppp-profile-egress

Syntax [**no**] **mlppp-profile-egress** *profile-id* [**create**]

Context config>qos

Description This command creates a profile for the user to configure the egress QoS parameters of a multiclass MLPPP bundle.

A maximum of 128 egress QoS profiles can be created on the system.

The **no** form of this command deletes the profile.

Parameters *profile-id* — Specifies an ingress mlppp-profile ID.

Values 1 — 65535

max-queue-size

Syntax **max-queue-size** *queue-size*
no max-queue-size

Context config>qos>mlppp-profile-egress>class

Description This command configures the maximum queue size for each MLPPP class queue for this profile.

Parameters *queue-size* — Specifies the maximum queue size.

Values 1 — 1000 in milliseconds of buffer space

mir

Syntax **mir** *mir-value*
no mir

Context config>qos>mlppp-profile-egress>class

Description This command configures the minimum information rate (MIR) scheduling parameter for each MLPPP class queue for this profile.

Parameters *mir-value* — Specifies the MIR scheduling parameter.

Values 1 — 100

weight

Syntax **weight** *weight-value*
no weight

Context config>qos>mlppp-profile-egress>class

Description This command configures the WRR weight scheduling parameter for each MLPPP class queue for this profile.

Parameters *weight-value* — Specifies the weight scheduling parameter.

Values 1 — 100

fc

Syntax **fc** *fc-name* **mlppp-class** *class-id*
no fc *fc-name*

Context config>qos>mlppp-profile-egress>class

Description This command configures the mapping of the system forwarding class to the MLPPP classes for this profile. There is a many-to-one relationship between the system forwarding class and an MLPPP class.

Parameters *fc-name* — Specifies the forwarding class name.

Values b3, l2, af, l1, h2, ef, h1, nc

class-id — Specifies the class ID.

Values 0 — 3

Standards and Protocol Support

Note that this Standards Compliance list is subject to change.

Ethernet Standards

IEEE 802.1ab-REV/D3 Station and Media Access Control Connectivity Discovery
IEEE 802.1d Bridging
IEEE 802.1p/Q VLAN Tagging
IEEE 802.1s Multiple Spanning Tree
IEEE 802.1w Rapid Spanning Tree Protocol
IEEE 802.1x Port Based Network Access Control
IEEE 802.1ad Provider Bridges
IEEE 802.1ah Provider Backbone Bridges
IEEE 802.1ag Service Layer OAM
IEEE 802.3ah Ethernet in the First Mile
IEEE 802.1ak Multiple MAC Registration Protocol
IEEE 802.3 10BaseT
IEEE 802.3ad Link Aggregation
IEEE 802.3ae 10Gbps Ethernet
IEEE 802.3ah Ethernet OAM
IEEE 802.3u 100BaseTX
IEEE 802.3x Flow Control
IEEE 802.3z 1000BaseSX/LX
ITU-T Y.1731 OAM functions and mechanisms for Ethernet based networks
ITU-T G.8031 Ethernet linear protection switching
ITU-T G.8032 Ethernet Ring Protection Switching (version 2)

OSPF

RFC 1765 OSPF Database Overflow
RFC 2328 OSPF Version 2
RFC 2370 Opaque LSA Support
RFC 2740 OSPF for IPv6 (OSPFv3)
RFC 3101 OSPF NSSA Option
RFC 3137 OSPF Stub Router Advertisement
RFC 3623 Graceful OSPF Restart – GR helper
RFC 3630 Traffic Engineering (TE) Extensions to OSPF Version 2
RFC 4203 OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS) - (support of Link Local/Remote Identifiers and SRLG sub-TLVs)
RFC 5185 OSPF Multi-Area Adjacency
RFC5243 OSPF Database Summary List Optimization

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1772 Application of BGP in the Internet
RFC 1965 Confederations for BGP
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MD5
RFC 2439 BGP Route Flap Dampening
RFC 2558 Multiprotocol Extensions for BGP-4
RFC 2918 Route Refresh Capability for BGP-4
RFC 3107 Carrying Label Information in BGP-4
RFC 3392 Capabilities Advertisement with BGP4
RFC 4271 BGP-4 (previously RFC 1771)
RFC 4360 BGP Extended Communities Attribute
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs) (previously RFC 2547bis BGP/MPLS VPNs)
RFC 4456 BGP Route Reflection: Alternative to Full-mesh IBGP
RFC 4486 Subcodes for BGP Cease Notification Message
RFC 4577 OSPF as the Provider/ Customer Edge Protocol for BGP/ MPLS IP Virtual Private Networks (VPNs)
RFC 4659 BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
RFC 4684 Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)
RFC 4724 Graceful Restart Mechanism for BGP – GR helper
RFC 4760 Multi-protocol Extensions for BGP
RFC 4798 Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)
RFC 4893 BGP Support for Four-octet AS Number Space
RFC 5004 Avoid BGP Best Path Transitions from One External to Another
RFC 5065 Confederations for BGP (obsoletes 3065)
RFC 5291 Outbound Route Filtering Capability for BGP-4
RFC 5575 Dissemination of Flow Specification Rules
RFC 5668 4-Octet AS Specific BGP Extended Community
draft-ietf-idr-add-paths Advertisement of Multiple Paths in BGP
Advertisement of the Best External Route in BGP
draft-ietf-idr-best-external

IS-IS

ISO/IEC 10589:2002, Second Edition Intermediate System to Intermediate System Intra-Domain Routeing Information Exchange Protocol
RFC 1195 Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
RFC 2973 IS-IS Mesh Groups
RFC 3359 Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System
RFC 3719 Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)
RFC 3787 Recommendations for Interoperable IP Networks using

Intermediate System to Intermediate System (IS-IS)
RFC 4971 Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information
RFC 5120 M-ISIS: Multi Topology (MT) Routing in IS-IS
RFC 5301 Dynamic Hostname Exchange Mechanism for IS-IS
RFC 5302 Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 5303 Three-Way Handshake for IS-IS Point-to-Point Adjacencies
RFC 5304 IS-IS Cryptographic Authentication
RFC 5305 IS-IS Extensions for Traffic Engineering TE
RFC 5306 Restart Signaling for IS-IS
RFC 5307 IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)
RFC 5309 Point-to-Point Operation over LAN in Link State Routing Protocols
RFC 5310 IS-IS Generic Cryptographic Authentication
RFC 6213 IS-IS BFD-Enabled TLV
RFC 6329 IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging
draft-ietf-isis-mi-02 IS-IS Multi-Instance

IPSec

RFC 2401 Security Architecture for the Internet Protocol
RFC 2406 IP Encapsulating Security Payload (ESP)
RFC 2409 The Internet Key Exchange (IKE)
RFC 2560 X.509 Internet Public Key Infrastructure
Online Certificate Status Protocol - OCSP
RFC 3706 IKE Dead Peer Detection
RFC 3947 Negotiation of NAT-Traversal in the IKE
RFC 3948 UDP Encapsulation of IPsec ESP Packets
RFC 4210 Internet X.509 Public Key Infrastructure
Certificate Management Protocol (CMP)
RFC 4211 Internet X.509 Public Key Infrastructure
Certificate Request Message Format (CRMF)
RFC 5996 Internet Key Exchange Protocol Version 2 (IKEv2)
RFC 5998 An Extension for EAP-Only Authentication in IKEv2
draft-ietf-ipsec-isakmp-xauth-06 – Extended Authentication within ISAKMP/Oakley (XAUTH)
draft-ietf-ipsec-isakmp-modecfg-05 – The ISAKMP Configuration Method

IPv6

RFC 1981 Path MTU Discovery for IPv6
RFC 2375 IPv6 Multicast Address Assignments
RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
RFC 2461 Neighbor Discovery for IPv6
RFC 2462 IPv6 Stateless Address Auto configuration
RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
RFC 2529 Transmission of IPv6 over IPv4 Domains without Explicit Tunnels
RFC 2545 Use of BGP-4 Multiprotocol Extension for IPv6 Inter-Domain Routing
RFC 2710 Multicast Listener Discovery (MLD) for IPv6
RFC 2740 OSPF for IPv6
RFC 3306 Unicast-Prefix-based IPv6 Multicast Addresses
RFC 3315 Dynamic Host Configuration Protocol for IPv6
RFC 3587 IPv6 Global Unicast Address Format
RFC3590 Source Address Selection for the Multicast Listener Discovery (MLD) Protocol
RFC 3810 Multicast Listener Discovery Version 2 (MLDv2) for IPv6
RFC 4007 IPv6 Scoped Address Architecture
RFC 4193 Unique Local IPv6 Unicast Addresses
RFC 4291 IPv6 Addressing Architecture
RFC 4443 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification
RFC 4552 Authentication/Confidentiality for OSPFv3
RFC 4659 BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
RFC 5072 IP Version 6 over PPP
RFC 5095 Deprecation of Type 0 Routing Headers in IPv6
draft-ietf-isis-ipv6-05
draft-ietf-isis-wg-multi-topology-xx.txt

Multicast

RFC 1112 Host Extensions for IP Multicasting (Snooping)
RFC 2236 Internet Group Management Protocol, (Snooping)
RFC 3376 Internet Group Management Protocol, Version 3 (Snooping)
RFC 2362 Protocol Independent Multicast-Sparse Mode (PIMSM)
RFC 3618 Multicast Source Discovery Protocol (MSDP)
RFC 3446 Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)
RFC 4601 Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)
RFC 4604 Using IGMPv3 and MLDv2 for Source-Specific Multicast
RFC 4607 Source-Specific Multicast for IP
RFC 4608 Source-Specific Protocol Independent Multicast in 232/8
RFC 4610 Anycast-RP Using Protocol Independent Multicast (PIM)
draft-ietf-pim-sm-bsr-06. Bootstrap Router (BSR) Mechanism for PIM
draft-rosen-vpn-mcast-15.txt Multicast in MPLS/BGP IP VPNs
draft-ietf-l3vpn-2547bis-mcast-07: Multicast in MPLS/BGP IP VPNs
draft-ietf-l3vpn-2547bis-mcast-bgp-05: BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs
RFC 3956: Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address

MPLS-GENERAL

RFC 2430 A Provider Architecture DiffServ & TE

RFC 2474 Definition of the DS Field the IPv4 and IPv6 Headers (Rev)

RFC 2597 Assured Forwarding PHB Group (rev3260)

RFC 2598 An Expedited Forwarding PHB

RFC 3031 MPLS Architecture

RFC 3032 MPLS Label Stack Encoding

RFC 3140 Per-Hop Behavior Identification Codes

RFC 3443 Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks

RFC 4182 Removing a Restriction on the use of MPLS Explicit NULL

RFC 4023 Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)

RFC 5332 MPLS Multicast Encapsulations

MPLS — LDP

RFC 3037 LDP Applicability

RFC 3478 Graceful Restart Mechanism for LDP – GR helper

RFC 5036 LDP Specification

RFC 5283 LDP extension for Inter-Area LSP

RFC 5443 LDP IGP Synchronization

RFC 6388 LDP Extensions for Point-to-Multipoint and Multipoint-to-Multipoint LSP

RFC 6826 Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

draft-pdutta-mpls-tldp-hello-reduce-04, Targeted LDP Hello Reduction

MPLS/RSVP-TE

RFC 2702 Requirements for Traffic Engineering over MPLS

RFC2747 RSVP Cryptographic Authentication

RFC 2961 RSVP Refresh Overhead Reduction Extensions

RFC3097 RSVP Cryptographic Authentication - Updated Message Type Value

RFC 3209 Extensions to RSVP for Tunnels

RFC 3473 Generalized Multi-Protocol Label Switching (GMPLS) Signaling

Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions – (support of IF_ID RSVP_HOP object with unnumbered interface and RSVP-TE Graceful Restart Helper Procedures)

RFC 3477 Signalling Unnumbered Links inResourceReSerVationProtocol-Traffic Engineering (RSVP-TE)

RFC 3564 Requirements for Diff-Serv-aware TE

RFC 3906 Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels

RFC 4090 Fast reroute Extensions to RSVP-TE for LSP Tunnels

RFC 4124 Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering

RFC 4125 Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering

RFC 4127 Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering

RFC 4561 Definition of a RRO Node-Id Sub-Object

RFC 4875 Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)

RFC 5151 Inter-domain MPLS and GMPLS Traffic Engineering – RSVP-TE Extensions

RFC 5712 MPLS Traffic Engineering Soft Preemption

RFC 5817 Graceful Shutdown in GMPLS Traffic Engineering Networks

draft-newton-mpls-te-dynamic-overbooking-00 A Diffserv-TE Implementation Model to dynamically change booking factors during failure events

MPLS — OAM

RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

RFC 6424 Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels

RFC 6425 Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping

MPLS-TP (7750/7450 only)

RFC 5586 MPLS Generic Associated Channel

RFC 5921 A Framework for MPLS in Transport Networks

RFC 5960 MPLS Transport Profile Data Plane Architecture

RFC 6370 MPLS-TP Identifiers

RFC 6378 MPLS-TP Linear Protection

RFC 6428 Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile

RFC 6426 MPLS On-Demand Connectivity and Route Tracing

RFC 6478 Pseudowire Status for Static Pseudowires

draft-ietf-mpls-tp-ethernet-addressing-02 MPLS-TP Next-Hop Ethernet Addressing

RIP

RFC 1058 RIP Version 1

RFC 2080 RIPng for IPv6

RFC 2082 RIP-2 MD5 Authentication

RFC 2453 RIP Version 2

TCP/IP

RFC 768 UDP

RFC 1350 The TFTP Protocol (Rev.

RFC 791 IP

RFC 792 ICMP

RFC 793 TCP

RFC 826 ARP

RFC 854 Telnet

RFC 951 BootP (rev)

RFC 1519 CIDR

RFC 1542 Clarifications and Extensions for the Bootstrap Protocol

RFC 1812 Requirements for IPv4 Routers

RFC 2347 TFTP option Extension
RFC 2328 TFTP Blocksize Option
RFC 2349 TFTP Timeout Interval and Transfer
Size option
RFC 2401 Security Architecture for Internet Protocol
RFC 2428 FTP Extensions for IPv6 and NATs
RFC 3596 DNS Extensions to Support IP version 6
RFC 5880 Bidirectional Forwarding Detection
RFC 5881 BFD IPv4 and IPv6 (Single Hop)
RFC 5883 BFD for Multihop Paths
RFC 5286 Basic Specification for IP Fast Reroute: Loop-Free Alternates
draft-litkowski-rtgwg-lfa-manageability-01 Operational management of Loop Free Alternates

VRRP

RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol
RFC 3768 Virtual Router Redundancy Protocol
draft-ietf-vrrp-unified-spec-02 Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6

PPP

RFC 1332 PPP IPCP
RFC 1377 PPP OSINLCP
RFC 1638/2878PPP BCP
RFC 1661 PPP (rev RFC2151)
RFC 1662 PPP in HDLC-like Framing
RFC 1877 PPP Internet Protocol Control Protocol Extensions for Name Server Addresses
RFC 1989 PPP Link Quality Monitoring
RFC 1990 The PPP Multilink Protocol (MP)
RFC 1994 PPP Challenge Handshake
Authentication Protocol (CHAP)
RFC 2516 A Method for Transmitting PPP Over Ethernet
RFC 2615 PPP over SONET/SDH
RFC 2686 The Multi-Class Extension to Multi-Link PPP

Frame Relay

FRF.1.2 - PVC User-to-Network Interface (UNI) Implementation Agreement
FRF.5 - Frame Relay/ATM PVC Network Interworking Implementation
ANSI T1.617 Annex D, DSS1 — Signalling Specification For Frame Relay Bearer Service.
FRF2.2. PVC Network-to- Network Interface (NNI) Implementation Agreement.
FRF.12 Frame Relay Fragmentation Implementation Agreement
FRF.16.1 Multilink Frame Relay UNI/NNI Implementation Agreement
ITU-T Q.933 Annex A Additional procedures for Permanent Virtual Connection (PVC) status management

ATM

RFC 1626 Default IP MTU for use over ATM AAL5
RFC 2514 Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management
RFC 2515 Definition of Managed Objects for ATM Management RFC 2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5
AF-TM-0121.000 Traffic Management Specification Version 4.1
ITU-T Recommendation I.610 B-ISDN Operation and Maintenance Principles and Functions version 11/95
ITU-T Recommendation I.432.1 BISDN user-network interface – Physical layer specification: General characteristics
GR-1248-CORE Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3
GR-1113-CORE Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1
AF-ILMI-0065.000 Integrated Local Management Interface (ILMI) Version 4.0
AF-TM-0150.00 Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR
AF-PHY-0086.001 Inverse Multiplexing for ATM (IMA) Specification Version 1.1

DHCP

RFC 2131 Dynamic Host Configuration Protocol (REV)
RFC 3046 DHCP Relay Agent Information Option (Option 82)
RFC 1534 Interoperation between DHCP and BOOTP

Policy Management and Credit Control

3GPP TS 29.212 Policy and Charging Control (PCC) over Gx/Sd Reference Point (Release 11) - Gx support as it applies to wireline environment (BNG)
RFC 3588 Diameter Base Protocol
RFC 4006 Diameter Credit Control Application

NAT

RFC 6333 Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion
RFC 6334 Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite
RFC 6888 Common Requirements For Carrier-Grade NATs (CGNs)
RFC 5508 NAT Behavioral Requirements for ICMP
RFC 5382 NAT Behavioral Requirements for TCP
RFC 6146 Statefull NAT64

VPLS

RFC 4761 Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling
RFC 4762 Virtual Private LAN Services Using LDP
RFC 5501 Requirements for Multicast Support in Virtual Private LAN Services
RFC 6074 Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)
draft-ietf-l2vpn-vpls-mcast-13. Multicast in VPLS

RFC 7041 Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging

Pseudowire

RFC 3985 Pseudo Wire Emulation Edge-to-Edge (PWE3)

RFC 4385 Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN

RFC 3916 Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)

RFC 4717 Encapsulation Methods for Transport ATM over MPLS Networks

RFC 4816 PWE3 ATM Transparent Cell Transport Service

RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks

RFC 4619 Encapsulation Methods for Transport of Frame Relay over MPLS Networks

RFC 4446 IANA Allocations for PWE3

RFC 4447 Pseudowire Setup and Maintenance Using LDP

RFC 5085 Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires

RFC 5659 An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge

RFC 5885 Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)

RFC 6310 Pseudowire (PW) OAM Message Mapping

RFC6391 Flow Aware Transport of Pseudowires over an MPLS PSN

RFC 6575 ARP Mediation for IP Interworking of Layer 2 VPN

RFC 6718 Pseudowire Redundancy

RFC 6870 Pseudowire Preferential Forwarding Status bit

draft-ietf-l2vpn-vpws-iw-oam-03 OAM Procedures for VPWS Interworking

draft-ietf-pwe3-mpls-eth-oam-iwk-07 MPLS and Ethernet OAM Interworking

draft-ietf-pwe3-dynamic-ms-pw-16 Dynamic Placement of Multi Segment Pseudo Wires

MFA Forum 9.0.0 The Use of Virtual trunks for ATM/MPLS Control Plane Interworking

MFA Forum 12.0.0 Multiservice Interworking - Ethernet over MPLS

MFA Forum 13.0.0 Fault Management for Multiservice Interworking v1.0

MFA Forum 16.0.0 Multiservice Interworking - IP over MPLS

ANCP/L2CP

RFC 5851 ANCP framework

draft-ietf-ancp-protocol-02 ANCP Protocol

Voice /Video Performance:

ITU-T G.107 The E Model- A computational model for use in planning.

ETSI TS 101 329-5 Annex E extensions- QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring

ITU-T Rec. P.564 Conformance testing for voice over IP transmission quality assessment models

ITU-T G.1020 Appendix I Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks- Mean Absolute Packet Delay Variation.& Markov Models.

RFC 3550 Appendix A.8- RTP A Transport Protocol for Real-Time Applications- Estimating the Interarrival Jitter.

Circuit Emulation

RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)

RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)

MEF-8 Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks, October 2004

RFC 5287 Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks

SONET/SDH

ITU-G.841 Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

AAA

RFC 2865 Remote Authentication Dial In User Service

RFC 2866 RADIUS Accounting

draft-grant-tacacs-02. The TACACS+ Protocol

SSH

RFC 4250 The Secure Shell (SSH) Protocol Assigned Numbers

RFC 4251 The Secure Shell (SSH) Protocol Architecture

RFC 4254 The Secure Shell (SSH) Connection Protocol

OpenFlow

ONF OpenFlow Switch Specification Version 1.3.1 (Hybrid-switch/ FlowTable)

Timing

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

ITU-T G.781 Telecommunication Standardization Section of ITU, Synchronization layer functions, issued 09/2008

ITU-T G.813 Telecommunication Standardization Section of ITU, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003.

GR-1244-CORE Clocks for the Synchronized Network: Common Generic Criteria, *Issue 3, May 2005*

ITU-T G.8261 Telecommunication Standardization Section of ITU, Timing and synchronization aspects in packet networks, issued 04/2008.

ITU-T G.8262 Telecommunication Standardization Section of ITU, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007.

ITU-T G.8264 Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008.

ITU-T G.8265.1 Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for frequency synchronization, issued 10/2010.

IEEE 1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

Network Management

ITU-T X.721 Information technology-OSI-Structure of Management Information

ITU-T X.734 Information technology-OSI-Systems Management: Event Report Management Function

M.3100/3120 Equipment and Connection Models

TMF 509/613 Network Connectivity Model

RFC 1157 SNMPv1

RFC 1215 A Convention for Defining Traps for use with the SNMP

RFC 1657 BGP4-MIB

RFC 1724 RIPv2-MIB

RFC 1850 OSPF-MIB

RFC 1907 SNMPv2-MIB

RFC 2011 IP-MIB

RFC 2138 RADIUS

RFC 2206 RSVP-MIB

RFC 2452 IPv6 Management Information Base for the

Transmission Control Protocol

RFC 2465 Management Information Base for IPv6: Textual Conventions and General Group

RFC 2558 SONET-MIB

RFC 2571 SNMP-FRAMEWORKMIB

RFC 2572 SNMP-MPD-MIB

RFC 2573 SNMP-TARGET-&-NOTIFICATION-MIB

RFC 2574 SNMP-USER-BASED-SMMIB

RFC 2575 SNMP-VIEW-BASEDACM-MIB

RFC 2576 SNMP-COMMUNITY-MIB

RFC 2578 Structure of Management Information Version 2 (SMIv2)

RFC 2665 EtherLike-MIB

RFC 2819 RMON-MIB

RFC 2863 IF-MIB

RFC 2864 INVERTED-STACK-MIB

RFC 2987 VRRP-MIB

RFC 3014 NOTIFICATION-LOGMIB

RFC 3019 IP Version 6 Management Information Base for The Multicast Listener Discovery Protocol

RFC 3164 Syslog

RFC 3273 HCRMON-MIB

RFC 3411 An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks

RFC 3412 Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)

RFC 3413 Simple Network Management Protocol (SNMP) Applications

RFC 3414 User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)

RFC 3418 SNMP MIB

RFC 3826 The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model

RFC 4113 Management Information Base for the User Datagram Protocol (UDP)

RFC 4292 IP-FORWARD-MIB

RFC 4293 MIB for the Internet Protocol

RFC 5101 Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information

RFC 6242 Using the NETCONF Protocol over Secure Shell (SSH)

draft-ietf-bfd-mib-00 Bidirectional Forwarding Detection Management Information Base

draft-ietf-isis-wg-mib-06 Management Information Base for Intermediate System to Intermediate System (IS-IS)

draft-ietf-ospf-mib-update-04 OSPF Version 2 Management Information Base

draft-ietf-mboned-msdp-mib-01 Multicast Source Discovery protocol MIB

draft-ietf-mpls-lsr-mib-06 Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base

draft-ietf-mpls-te-mib-04 Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base

draft-ietf-mpls-ldp-mib-07 MPLS Label Switch Router Management Information Base Using SMIv2

IANA ifType MIB

IEEE 802.3- LAG-MIB