



**7450 ETHERNET SERVICE SWITCH
7750 SERVICE ROUTER
7950 EXTENSIBLE ROUTING SYSTEM
VIRTUALIZED SERVICE ROUTER**

**QUALITY OF SERVICE GUIDE
RELEASE 16.0.R4**

3HE 14134 AAAB TQZZA 02

Issue: 01

November 2018

Nokia is a registered trademark of Nokia Corporation. Other products and company names mentioned herein may be trademarks or tradenames of their respective owners.

The information presented is subject to change without notice. No responsibility is assumed for inaccuracies contained herein.

© 2018 Nokia.

Contains proprietary/trade secret information which is the property of Nokia and must not be made available to, or copied or used by anyone outside Nokia without its written authorization. Not to be used or disclosed except in accordance with applicable agreements.

Table of Contents

1	Getting Started	17
1.1	About This Guide	17
1.2	QoS Configuration Process	19
2	QoS Policies	21
2.1	QoS Overview	21
2.2	Forwarding Classes	22
2.2.1	High-Priority Classes	23
2.2.2	Assured Classes	23
2.2.3	Best-Effort Classes	24
2.3	Queue Parameters	25
2.3.1	Queue ID	25
2.3.2	Unicast or Multipoint Queue	25
2.3.3	Queue Type	25
2.3.4	Queue Scheduling	26
2.3.5	Peak Information Rate	27
2.3.6	Committed Information Rate	28
2.3.7	Fair Information Rate	29
2.3.8	Adaptation Rule	30
2.3.9	Committed Burst Size	31
2.3.10	Maximum Burst Size	32
2.3.11	Queue Drop Tails	32
2.3.12	WRED Per Queue	34
2.3.12.1	Native Queue Mode	35
2.3.12.2	Pool Per Queue Mode	38
2.3.13	Packet Markings	42
2.3.14	Queue Counters	43
2.3.15	Color Aware Profiling	43
2.4	QoS Policies Overview	45
2.4.1	Service versus Network QoS	48
2.4.2	QoS Policy Entities	49
2.4.3	Network QoS Policies	50
2.4.4	Network Queue QoS Policies	52
2.4.5	Service Ingress QoS Policies	53
2.4.5.1	FC Mapping Based on EXP Bits at VLL/VPLS SAP	55
2.4.5.2	Egress Forwarding Class Override	59
2.4.6	Service Egress QoS Policies	60
2.4.7	Named Pool Policies	62
2.4.8	Slope Policies	65
2.4.8.1	RED Slopes	65
2.4.8.2	Tuning the SBAU Calculation	68
2.4.8.3	Slope Policy Parameters	70
2.4.9	Scheduler Policies	71
2.4.9.1	Virtual Hierarchical Scheduling	74
2.4.9.2	Hierarchical Scheduler Policies	74

2.4.9.3	Tiers.....	75
2.4.9.4	Scheduler Policies Applied to Applications.....	76
2.4.9.5	Scheduler Policies Applied to SAPs	77
2.4.9.6	Scheduler Policies Applied to Customer SLAs.....	78
2.4.9.7	Scheduler Policies Applied to Multiservice Sites	78
2.4.10	ATM Traffic Descriptor Profiles.....	78
2.4.11	Configuration Notes.....	79
3	Match List for QoS Policies	81
3.1	Match List for QoS Policies Commands.....	83
3.1.1	Command Hierarchies.....	83
3.1.1.1	Configuration Commands.....	83
3.1.1.2	Show Commands	83
3.1.2	Command Descriptions	84
3.1.2.1	Configuration Commands.....	84
3.1.2.2	Show Commands	86
4	Network QoS Policies.....	91
4.1	Network QoS Policies Overview	91
4.1.1	Storing Match-Criteria Entries	92
4.1.1.1	FP4-based Cards	92
4.2	Network Ingress.....	94
4.2.1	Network Ingress Tunnel QoS Override.....	94
4.2.2	Network Ingress IP Match Criteria.....	95
4.2.3	Network Ingress IPv6 Match Criteria	96
4.3	Network Egress	98
4.3.1	Egress Packet Reclassification Based on IP precedence DSCP	99
4.3.2	Network Egress IP Match Criteria	100
4.3.3	Network Egress IPv6 Match Criteria.....	100
4.4	QoS for Self-Generated (CPU) Traffic on Network Interfaces.....	102
4.4.1	Default DSCP Mapping Table	105
4.5	Basic Configurations.....	106
4.5.1	Creating a Network QoS Policy	106
4.5.2	Applying Network QoS Policies	110
4.5.3	Default Network QoS Policy Values	110
4.6	Service Management Tasks	116
4.6.1	Deleting QoS Policies.....	116
4.6.2	Removing a Policy from the QoS Configuration	117
4.6.3	Copying and Overwriting Network Policies.....	117
4.6.4	Editing QoS Policies	118
4.7	Network QoS Policy Command Reference	119
4.7.1	Command Hierarchies.....	119
4.7.1.1	Configuration Commands.....	119
4.7.1.2	Self-Generated Traffic Commands.....	123
4.7.1.3	Operational Commands.....	123
4.7.1.4	Show Commands	124
4.7.1.5	Debug Commands.....	124
4.7.2	Command Descriptions	124
4.7.2.1	Configuration Commands.....	125

4.7.2.2	Self-Generated Traffic Commands	162
4.7.2.3	Operational Commands	164
4.7.2.4	Show Commands	165
5	Network Queue QoS Policies	177
5.1	Overview	177
5.2	Network Queue Parent Scheduler	178
5.3	Basic Configurations	179
5.3.1	Creating a Network Queue QoS Policy	179
5.3.2	Applying Network Queue QoS Policies	179
5.3.2.1	XMASs and MDAs	179
5.3.2.2	Ethernet Ports	180
5.3.2.3	SONET/SDH Ports	181
5.3.3	Default Network Queue Policy Values	181
5.4	Service Management Tasks	190
5.4.1	Deleting QoS Policies	190
5.4.2	Removing a Policy from the QoS Configuration	190
5.4.3	Copying and Overwriting QoS Policies	190
5.4.4	Editing QoS Policies	191
5.5	Network Queue QoS Policy Command Reference	193
5.5.1	Command Hierarchies	193
5.5.1.1	Configuration Commands	193
5.5.1.2	Operational Commands	194
5.5.1.3	Show Commands	194
5.5.2	Command Descriptions	195
5.5.2.1	Configuration Commands	195
5.5.2.2	Operational Commands	218
5.5.2.3	Show Commands	219
6	Service Ingress and Egress QoS Policies	227
6.1	Overview	227
6.2	Basic Configurations	228
6.3	Service Ingress QoS Policy	229
6.3.1	Service Ingress QoS Queue	229
6.3.2	Ingress Percent-Rate Support	231
6.3.3	Ingress Forwarding Class (FC)	232
6.3.4	Ingress IP Match Criteria	233
6.3.5	Ingress IPv6 Match Criteria	233
6.3.6	Ingress MAC Match Criteria	234
6.3.6.1	VLAN ID Matching	235
6.3.7	Ingress Criteria Classification Directly to Policer	236
6.3.8	Virtual Network Identifier (VNI) Classification	237
6.3.9	FC Mapping Based on EXP Bits	238
6.3.10	Storing Match Criteria Entries	239
6.3.10.1	FP4-based Cards	239
6.4	Service Egress QoS Policy	241
6.4.1	Service Egress QoS Queue	241
6.4.2	Egress Percent-Rate Support	242
6.4.3	Dynamic MBS for Egress Queue Group Queues	242

6.4.3.1	Queue Length as a Delay Value.....	245
6.4.4	Egress SAP FC and FP Overrides	246
6.4.5	Egress IP Match Criteria.....	247
6.4.6	Egress IPv6 Match Criteria.....	247
6.4.7	Egress Criteria Classification Directly to Policer.....	248
6.4.8	Storing Match Criteria Entries.....	250
6.4.8.1	FP4-based Cards	250
6.4.9	Dot1p Egress Remarking	251
6.4.9.1	DEI Egress Remarking	252
6.4.10	DSCP and IP Precedence Egress Remarking	257
6.4.11	Queue Depth Monitoring	257
6.5	Service Management Tasks	260
6.5.1	Applying Service Ingress and Egress Policies	260
6.5.1.1	Epipe	260
6.5.1.2	IES.....	261
6.5.1.3	VPLS	261
6.5.1.4	VPRN.....	262
6.5.2	Editing QoS Policies	262
6.5.3	Copying and Overwriting QoS Policies.....	262
6.5.4	Deleting QoS Policies.....	263
6.5.5	Removing a Policy from the QoS Configuration	264
6.6	Service Ingress and Egress QoS Policy Command Reference.....	265
6.6.1	Command Hierarchies.....	265
6.6.1.1	Service Ingress QoS Policy Commands.....	265
6.6.1.2	Service Egress QoS Policy Commands	269
6.6.1.3	Operational Commands.....	273
6.6.1.4	Show Commands	274
6.6.1.5	Debug Commands.....	274
6.6.2	Command Descriptions	274
6.6.2.1	Generic Commands.....	275
6.6.2.2	Service Ingress QoS Policy Commands.....	275
6.6.2.3	Service Egress QoS Policy Commands	362
6.6.2.4	Operational Commands.....	440
6.6.2.5	Show Commands	441
6.6.2.6	Debug Commands.....	458
7	Queue Sharing and Redirection	461
7.1	Queue Sharing and Redirection	461
7.1.1	Supported Platforms.....	461
7.2	Queue Group Applications	462
7.2.1	Access SAP Queue Group Applications.....	462
7.2.1.1	Ingress Per SAP Statistics with Ingress Queue Groups.....	463
7.2.1.2	Ingress Access Port Queue Group Hardware Queue Allocation.....	465
7.2.2	Network Port Queue Groups for IP Interfaces.....	465
7.2.3	Pseudowire Shaping for Layer 2 and Layer 3 Services.....	466
7.2.4	QoS on Ingress Bindings.....	467
7.2.5	VXLAN VNI Queue Group Redirection.....	468
7.2.5.1	Queue Group Redirect List.....	469
7.2.5.2	Queue Group Redirect List Example.....	472

7.3	Queue Group Templates	475
7.4	Port Queue Groups	476
7.4.1	Percent-Rate Support.....	476
7.5	Forwarding Plane Queue Groups.....	477
7.6	Redirection Models.....	478
7.7	Access SAP Forwarding Class-Based Redirection	479
7.7.1	Ingress and Egress SAP Forwarding Class Redirection Association Rules	481
7.7.1.1	Policy Based Provisioning Model	481
7.7.1.2	SAP-Based Provisioning Model.....	484
7.7.2	Access Queue Group Statistics.....	485
7.7.2.1	Port Queue Groups	485
7.7.2.2	Forwarding Plane Queue Groups.....	486
7.8	Network IP Interface Forwarding Class-Based Redirection	487
7.8.1	Egress Network Forwarding Class Redirection Association Rules	487
7.8.2	Egress Network IP Interface Statistics	489
7.9	PW Shaping.....	490
7.9.1	Ingress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection	490
7.9.1.1	Feature Configuration.....	490
7.9.1.2	Provisioning Model	491
7.9.1.3	Ingress Packet Classification.....	493
7.9.2	Egress PW Shaping using Spoke-SDP Forwarding Class-Based Redirection	494
7.9.2.1	Feature Configuration.....	494
7.9.2.2	Provisioning Model	495
7.9.2.3	Egress Marking of PW Packet Header	497
7.9.2.4	Ingress and Egress PW Statistics	497
7.10	Queue Group Behavior on LAG	499
7.10.1	Queue Group Queue Instantiation Per Link	499
7.10.2	Per-Link Queue Group Queue Parameters	499
7.10.3	Adding a Queue Group to an Existing LAG.....	499
7.10.4	Adding a Port to a LAG.....	499
7.10.5	Removing a Queue Group from a LAG	500
7.11	Basic Configurations.....	501
7.11.1	Configuring an Ingress Queue Group Template.....	501
7.11.2	Configuring Egress Queue Group Template	501
7.11.3	Applying Ingress Queue Group to SAP Ingress Policy.....	502
7.11.4	Applying Egress Queue Group to SAP Egress Policy.....	503
7.11.5	Configuring SAP-based Egress Queue Redirection	503
7.11.6	Configuring Queue Group on Ethernet Access Ingress Port.....	505
7.11.7	Configuring Overrides.....	506
7.11.8	Configuring Queue Group on Ethernet Access Egress Port	507
7.11.9	Configuring Queue Group for Network Egress Traffic on Port	508
7.11.10	Configuring Queue Group for Network Ingress Traffic on Forwarding Plane	509
7.11.11	Using Queue Groups to Police Ingress/Egress Traffic on Network Interface.....	510

7.11.12	Configuring Ingress/Egress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection	511
7.11.13	Specifying QoS Policies on Service SAPs	514
7.12	QoS Queue Group Template Command Reference	515
7.12.1	Command Hierarchies	515
7.12.1.1	Configuring Egress Queue Group Templates	515
7.12.1.2	Configuring Ingress Queue Group Templates	517
7.12.1.3	Configuring VXLAN VNI Queue Group Redirection	518
7.12.1.4	Show Commands	519
7.12.1.5	Monitor Commands	519
7.12.2	Command Descriptions	520
7.12.2.1	Configuration Commands	520
7.12.2.2	Show Commands	586
7.12.2.3	Monitor Commands	596
8	Scheduler QoS Policies	601
8.1	Scheduler Policies	601
8.1.1	Egress Port-Based Schedulers	601
8.1.1.1	Service/Subscriber or Multiservice Site Egress Port Bandwidth Allocation	604
8.1.1.2	Service or Subscriber or Multiservice Site Scheduler Child to Port Scheduler Parent	606
8.1.1.3	Frame and Packet-Based Bandwidth Allocation	609
8.1.1.4	Parental Association Scope	612
8.1.1.5	Service or Subscriber or Multiservice Site-Level Scheduler Parental Association Scope	612
8.1.1.6	Network Queue Parent Scheduler	613
8.1.1.7	Foster Parent Behavior for Orphaned Queues and Schedulers	614
8.1.1.8	Congestion Monitoring on Egress Port Scheduler	615
8.1.2	Frame-Based Accounting	620
8.1.2.1	Operational Modifications	621
8.1.2.2	Existing Egress Port-Based Virtual Scheduling	621
8.1.2.3	Behavior Modifications for Frame-Based Accounting	621
8.1.2.4	Virtual Scheduler Rate and Queue Rate Parameter Interpretation	622
8.1.3	Virtual Scheduling Unused Bandwidth Distribution	623
8.1.3.1	Default Unused Bandwidth Distribution	623
8.1.3.2	Limit Unused Bandwidth	624
8.1.4	Configuring Port Scheduler Policies	626
8.1.4.1	Port Scheduler Structure	626
8.1.4.2	Special Orphan Queue and Scheduler Behavior	626
8.1.4.3	Packet to Frame Bandwidth Conversion	627
8.1.4.4	Aggregate Rate Limits for Directly Attached Queues	628
8.1.4.5	SAP Egress QoS Policy Queue Parenting	629
8.1.4.6	Network Queue QoS Policy Queue Parenting	629
8.1.4.7	Egress Port Scheduler Overrides	629
8.1.4.8	Applying a Port Scheduler Policy to a Virtual Port	629
8.1.4.9	Applying Aggregate Rate Limit to a Vport	631
8.1.4.10	Applying a Scheduler Policy to a Vport	632
8.1.4.11	Weighted Scheduler Group in a Port Scheduler Policy	633

8.2	Basic Configurations	634
8.2.1	Creating a QoS Scheduler Policy	634
8.2.2	Applying Scheduler Policies	636
8.2.2.1	Customer	636
8.2.2.2	Epipe	636
8.2.2.3	IES	637
8.2.2.4	VPLS	638
8.2.2.5	VPRN	639
8.2.3	Creating a QoS Port Scheduler Policy	639
8.2.4	Configuring Port Parent Parameters	640
8.2.4.1	Within-CIR Priority Level Parameters	641
8.2.4.2	Above-CIR Priority Level Parameters	641
8.2.5	Configuring Distributed LAG Rate	642
8.3	Service Management Tasks	655
8.3.1	Deleting QoS Policies	655
8.3.1.1	Removing a QoS Policy from a Customer Multiservice Site	655
8.3.1.2	Removing a QoS Policy from SAP(s)	655
8.3.1.3	Removing a Policy from the QoS Configuration	656
8.3.2	Copying and Overwriting Scheduler Policies	657
8.3.3	Editing QoS Policies	658
8.4	QoS Scheduler Policy Command Reference	659
8.4.1	Command Hierarchies	659
8.4.1.1	Scheduler Policy Configuration Commands	659
8.4.1.2	Port Scheduler Policy Configuration Commands	659
8.4.1.3	Operational Commands	660
8.4.1.4	Show Commands	660
8.4.1.5	Clear Commands	662
8.4.2	Command Descriptions	662
8.4.2.1	Configuration Commands	662
8.4.2.2	Scheduler Policy Commands	663
8.4.2.3	Port Scheduler Policy Commands	671
8.4.2.4	Operational Commands	679
8.4.2.5	Show Commands	680
8.4.2.6	Clear Commands	718
9	Slope QoS Policies	723
9.1	Overview	723
9.2	Basic Configurations	724
9.2.1	Creating a Slope QoS Policy	724
9.2.2	Applying Slope Policies	725
9.2.2.1	Ports	725
9.2.2.2	FP	726
9.2.3	Default Slope Policy Values	726
9.3	Service Management Tasks	729
9.3.1	Deleting QoS Policies	729
9.3.1.1	Ports	729
9.3.1.2	FP	729
9.3.1.3	Removing a Policy from the QoS Configuration	730
9.3.2	Copying and Overwriting QoS Policies	730

9.3.3	Editing QoS Policies	730
9.4	Slope QoS Policy Command Reference	731
9.4.1	Command Hierarchies	731
9.4.1.1	Configuration Commands	731
9.4.1.2	Operational Commands	732
9.4.1.3	Show Commands	732
9.4.2	Command Descriptions	732
9.4.2.1	Configuration Commands	732
9.4.2.2	Slope Policy QoS Commands	733
9.4.2.3	Operational Commands	738
9.4.2.4	Show Commands	738
10	Shared-Queue QoS Policies	743
10.1	Overview	743
10.2	Multipoint Shared Queuing	744
10.2.1	Ingress Queuing Modes of Operation	744
10.2.2	Ingress Service Queuing	744
10.2.2.1	Ingress Shared Queuing	745
10.2.2.2	Ingress Multipoint Shared Queuing	747
10.3	Basic Configurations	750
10.3.1	Modifying the Default Shared-Queue Policy	750
10.3.1.1	Applying Shared-Queue Policies	751
10.3.2	Default Shared Queue Policy Values	753
10.4	Shared-Queue QoS Policy Command Reference	761
10.4.1	Command Hierarchies	761
10.4.1.1	Configuration Commands	761
10.4.1.2	Show Commands	761
10.4.2	Command Descriptions	762
10.4.2.1	Configuration Commands	762
10.4.2.2	Show Commands	771
11	Advanced QoS Policies	773
11.1	Overview	773
11.2	Basic Configurations	774
11.2.1	Creating an Advanced QoS Policy	774
11.2.2	Applying Advanced Policies	775
11.2.2.1	Queue Group	775
11.2.2.2	SAP Ingress	775
11.2.2.3	SAP Egress	775
11.2.3	Default Advanced Policy Values	775
11.3	Service Management Tasks	777
11.3.1	Deleting QoS Policies	777
11.3.1.1	Queue Group	777
11.3.1.2	SAP Ingress	777
11.3.1.3	SAP Egress	777
11.3.2	Copying and Overwriting Advanced Policies	778
11.3.3	Editing Advanced Policies	778
11.4	Advanced QoS Policy Command Reference	779
11.4.1	Command Hierarchies	779

11.4.1.1	Configuration Commands.....	779
11.4.1.2	Operational Commands.....	780
11.4.1.3	Show Commands.....	780
11.4.2	Command Descriptions.....	780
11.4.2.1	Configuration Commands.....	780
11.4.2.2	Advanced Policy QoS Commands.....	781
11.4.2.3	Operational Commands.....	793
11.4.2.4	Show Commands.....	794
12	QoS ATM Traffic Descriptor Profiles	799
12.1	ATM Traffic Management.....	799
12.2	QoS Model for ATM-Based Services.....	800
12.3	ATM Service Categories.....	802
12.4	ATM Traffic Descriptors and QoS Parameters.....	803
12.5	Policing.....	804
12.6	Shaping.....	805
12.7	ATM Queuing and Scheduling.....	806
12.8	Congestion Avoidance.....	807
12.9	Basic Configurations.....	808
12.9.1	Creating an ATM-TD-Profile QoS Policy.....	808
12.9.2	Applying ATM-TD-Profile Policies.....	808
12.9.2.1	ATM VLL (Apipe) SAPs.....	809
12.9.2.2	Epipe SAPs.....	809
12.9.2.3	IES SAPs.....	809
12.9.2.4	Lpipe SAPs.....	809
12.9.2.5	VPRN SAPs.....	810
12.9.2.6	VPLS SAPs.....	810
12.9.3	Default ATM-TD-Profile Policy Values.....	810
12.10	Service Management Tasks.....	812
12.10.1	Removing a Profile from the QoS Configuration.....	812
12.10.2	Copying and Overwriting Profiles.....	812
12.10.3	Editing QoS Policies.....	812
12.11	ATM QoS Policy Command Reference.....	813
12.11.1	Command Hierarchies.....	813
12.11.1.1	Configuration Commands.....	813
12.11.1.2	Operational Commands.....	813
12.11.1.3	Show Commands.....	813
12.11.2	Command Descriptions.....	814
12.11.2.1	Configuration Commands.....	814
13	Named Pools	829
13.1	Overview.....	829
13.2	Named Pool Mode for IOM3-XP Card.....	833
13.3	Basic Configurations.....	834
13.3.1	Creating a Named Pool QoS Policy.....	834
13.3.1.1	Named Pool Configuration Procedure.....	834
13.3.2	Allocation Steps.....	836
13.4	Named Pools QoS Policy Command Reference.....	839
13.4.1	Command Hierarchies.....	839

13.4.1.1	Configuration Commands.....	839
13.4.1.2	Show Commands.....	839
13.4.2	Command Descriptions.....	840
13.4.2.1	Configuration Commands.....	840
13.4.2.2	Named Pool Policy Creation.....	840
13.4.2.3	Show Commands.....	849
14	High Scale Ethernet MDA Capabilities	853
14.1	HSMDA QoS Model.....	853
14.1.1	Queue Scaling.....	854
14.1.2	Port-Based Scheduling.....	854
14.1.3	Dual Pass Queuing.....	858
14.1.4	Egress Intermediate Destination Secondary Shapers.....	858
14.1.5	Packet and Octet Counting.....	860
14.1.6	Above-CIR Discard with PIR Bypass.....	862
14.1.6.1	HSMDA Ingress Queue Policing Mode	862
14.1.7	HSMDA Buffer Utilization Controls	864
14.1.7.1	HSMDA Buffer Pools	864
14.1.7.2	Identifying Queue Groups as Provisioned or System.....	864
14.1.7.3	Provisioned and System Port Class Pools	865
14.1.7.4	Aggregate Pools for Type and Class Separation	866
14.1.7.5	Use of Aggregate Control Buffer Pools	867
14.1.7.6	HSMDA Buffer Pool Policy	870
14.1.7.7	Port Class Pool Sizing	872
14.1.7.8	HSMDA Available Buffer Register Operation	872
14.1.8	HSMDA Queue Congestion and Buffer Utilization Controls	873
14.1.8.1	Maximum HSMDA Queue Depth.....	874
14.1.8.2	Control Plane HSMDA RED Slope Policy Management	874
14.1.8.3	HSMDA Slope Policy MBS Parameter	875
14.1.8.4	HSMDA Slope Policy Slope Parameters	875
14.1.8.5	HSMDA Slope Shutdown Behavior	879
14.1.8.6	Ingress Packet Mapping to HSMDA RED Slope	879
14.1.8.7	Egress Packet Mapping to HSMDA RED Slope.....	879
14.1.8.8	HSMDA Queue Congestion or Pool Congestion Discard Stats.....	880
14.1.9	Egress Queue CIR-Based Dot1p Remarking.....	881
14.2	SAP Ingress and SAP Egress QoS Policies.....	882
14.2.1	SAP Ingress QoS Policy.....	882
14.2.2	SAP Egress QoS Policy	883
14.3	Subscriber Queuing Differences.....	884
14.4	Basic HSMDA Configurations.....	885
14.4.1	HSMDA Pool Policies	885
14.4.2	HSMDA Scheduler Policies	885
14.4.3	HSMDA Slope Policies	886
14.4.4	Egress Queue Group.....	887
14.4.5	Configuring HSMDA Queue Group Overrides	887
14.5	Applying HSMDA Policies	889
14.6	HSMDA Features	890
14.6.1	HSMDA LAG	890
14.6.2	Billing.....	890

14.6.3	Resource Management	891
14.6.4	HSMDA Queue Groups	891
14.6.4.1	Scheduling Classes	892
14.6.4.2	Scheduling Class Weighted Groups	892
14.6.4.3	Scheduler Strict Priority Levels	892
14.6.4.4	Strict Priority Level PIR	893
14.6.4.5	Scheduler Maximum Rate	893
14.6.4.6	HSMDA Scheduler Policy Overrides	893
14.6.4.7	Orphan Queues	894
14.6.4.8	Default HSMDA Scheduling Policy	894
14.7	HSMDA Command Reference	895
14.7.1	Command Hierarchies	895
14.7.1.1	Show Commands	896
14.7.2	Command Descriptions	896
14.7.2.1	Configuration Commands	896
15	High Scale QoS IOM	931
15.1	Introduction	931
15.1.1	HSQ Shaping	931
15.1.2	HSQ Scheduling	933
15.1.3	HSQ Buffer Management	934
15.1.4	HSQ LAGs	935
15.2	HSQ IOM-Specific Configurations	936
15.3	Non-Applicable HSQ IOM Functions	937
15.4	High Scale QoS IOM Command Reference	939
15.4.1	Command Hierarchies	939
15.4.1.1	HS Attachment Policy Commands	939
15.4.1.2	HS Pool Policy Commands	940
15.4.1.3	HS Port Pool Policy Commands	940
15.4.1.4	HS Scheduler Policy Commands	941
15.4.1.5	Network Queue QoS Policy HS Commands	941
15.4.1.6	SAP Egress Policy QoS HS Commands	942
15.4.1.7	Queue Group Template HS Commands	942
15.4.1.8	Services Configuration HS Commands	943
15.4.1.9	Forwarding Plane HS Commands	946
15.4.1.10	Port Configuration HS Commands	947
15.4.1.11	LAG HS Commands	947
15.4.1.12	Operational Commands	948
15.4.1.13	Show Commands	948
15.4.1.14	Clear Commands	948
15.4.2	Command Descriptions	950
15.4.2.1	Configuration Commands	950
15.4.2.2	HS Attachment Policy Commands	950
15.4.2.3	HS Pool Policy Commands	956
15.4.2.4	HS Port Pool Policy Commands	965
15.4.2.5	HS Scheduler Policy Commands	972
15.4.2.6	Network Queue HS Attachment Policy Commands	977
15.4.2.7	SAP Egress HS Attachment Policy Commands	982
15.4.2.8	Queue Group Template Policy HS Commands	986

15.4.2.9	Operational Commands.....	991
15.4.2.10	Show Commands	992
15.4.2.11	Clear Commands.....	1007
16	QoS in MC-MLPPP	1009
16.1	Overview.....	1009
16.2	Ingress MLPPP Class Reassembly	1013
16.3	Basic Configurations.....	1014
16.3.1	Configuring MC-MLPPP	1015
16.3.2	Configuring MC-MLPPP QoS Parameters	1015
16.4	QoS in MLFR and FRF.12 Fragmentation.....	1018
16.4.1	QoS in MLFR.....	1018
16.4.2	QoS in FRF.12 End-to-End Fragmentation	1020
16.5	MLPPP Command Reference	1023
16.5.1	Command Hierarchies.....	1023
16.5.1.1	Configuration Commands.....	1023
16.5.2	Command Descriptions	1024
16.5.2.1	Configuration Commands.....	1024
17	Class Fair Hierarchical Policing (CFHP).....	1029
17.1	Overview.....	1029
17.2	Parent Policer Priority and Unfair Sensitive Discard Thresholds.....	1031
17.3	CFHP Ingress and Egress Use Cases	1033
17.4	Post-CFHP Queuing and Scheduling	1034
17.4.1	Ingress CFHP Queuing	1034
17.4.2	Egress CFHP Queuing.....	1035
17.4.2.1	Policer to Local Queue Mapping	1035
17.4.3	Egress Subscriber CFHP Queuing.....	1036
17.4.3.1	Subscriber Destination String Queue Group Identification	1037
17.4.4	SAP Default Destination String.....	1038
17.5	CFHP Policer Control Policy.....	1039
17.5.1	Policer Control Policy Root Arbiter	1039
17.5.2	Tier 1 and Tier 2 Explicit Arbiters	1040
17.5.3	Explicit Arbiter Rate Limits.....	1040
17.5.4	CFHP with Child Policer Exceed PIR Enabled	1040
17.6	CFHP Child Policer Definition and Creation	1042
17.7	Policer Enabled SAP QoS Policy Applicability	1043
17.8	Child Policer Parent Association	1044
17.9	Profile-Capped Policers.....	1045
17.10	Policer Interaction with Profile, Discard Eligibility, and Ingress Priority	1049
17.10.1	Ingress 'Undefined' Initial Profile	1051
17.10.2	Ingress Explicitly 'In-Profile' State Packet Handling without Profile- Capped Mode	1051
17.10.3	Ingress Explicitly 'In-Profile' State Packet Handling with Profile- Capped Mode	1052
17.10.4	Ingress Explicit 'Out-of-Profile' State Packet Handling	1052
17.10.5	Egress Explicit Profile Reclassification.....	1053
17.10.6	Preserving Out of Profile State at Egress Policer.....	1053

17.10.7	Egress Policer CIR Packet Handling without Profile-capped Mode	1053
17.10.8	Egress Policer CIR Packet Handling with Profile-capped Mode	1054
17.10.9	Forwarding Traffic Exceeding PIR in Egress Policers	1054
17.10.10	Post Egress Policer Packet Forwarding Class and Profile State Remapping	1055
17.10.11	Ingress Child Policer Stat-Mode	1057
17.10.12	Egress Child Policer Stat-Mode.....	1060
17.11	Profile-Preferred Mode Root Policers	1062
17.12	Child Policer Hierarchical QoS Parenting.....	1064
17.13	Class Fair Hierarchical Policing (CFHP) Policy Command Reference	1069
17.13.1	Command Hierarchies	1069
17.13.1.1	Class Fair Hierarchical Policing Commands	1069
17.13.1.2	Operational Commands.....	1069
17.13.1.3	Show Commands	1070
17.13.2	Command Descriptions	1070
17.13.2.1	Configuration Commands.....	1070
18	Frequently Used QoS Terms	1087
18.1	Overview.....	1087
18.2	Above-CIR Distribution	1088
18.3	Available Bandwidth	1089
18.4	CBS	1090
18.5	CIR	1091
18.6	CIR Level.....	1092
18.7	CIR Weight	1093
18.8	Child	1094
18.9	Level	1095
18.10	MBS.....	1096
18.11	Offered Load.....	1097
18.12	Orphan.....	1098
18.13	Parent.....	1099
18.14	Queue.....	1100
18.15	Rate	1101
18.16	Root (Scheduler)	1102
18.17	Scheduler Policy.....	1103
18.18	Tier	1104
18.19	Virtual Scheduler	1105
18.20	Weight	1106
18.21	Within-CIR Distribution	1107
19	Standards and Protocol Support	1109

1 Getting Started

1.1 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.

The topics and commands described in this document apply to the:

- 7450 ESS
- 7750 SR
- 7950 XRS
- VSR

[Table 1](#) lists the available chassis types for each SR OS router.

Table 1 Supported SR OS Router Chassis Types

7450 ESS	7750 SR	7950 XRS
• 7450 ESS-7/12 running in standard mode (not mixed-mode)	• 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e • 7750 SR-7s/14-s • 7750 SR-1	• 7950 XRS-16c • 7950 XRS-20/40

For a list of unsupported features by platform and chassis, refer to the SR OS 16.0.Rx Software Release Notes, part number 3HE 14220 000x TQZZA or the VSR Release Notes, part number 3HE 14204 000x TQZZA.

Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.



Note: This guide generically covers Release 16.0.Rx content and may contain some content that will be released in later maintenance loads. Refer to the SR OS 16.0.Rx Software Release Notes, part number 3HE 14220 000x TQZZA or the *VSR Release Notes*, part number 3HE 14204 000x TQZZA, for information about features supported in each load of the Release 16.0.Rx software.

1.2 QoS Configuration Process

[Table 2](#) 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 2 Configuration Process

Area	Task	Section
Match lists in network, SAP ingress and SAP egress QoS policies	Configure a QoS match list	Match List for QoS Policies
Network QoS policies	Configure a network QoS policy	Basic Configurations
	Service management	Service Management Tasks
Network queue QoS policies	Configure a network queue QoS policy	Basic Configurations
	Service management	Service Management Tasks
Service Ingress and Egress QoS policies	Configure a service ingress QoS policy	Service Ingress QoS Policy
	Configure a service egress QoS policy	Service Egress QoS Policy
	Service management	Service Management Tasks
Queue sharing and redirection	Configure ingress and egress queue groups	Basic Configurations
Scheduler QoS policies	Configure a QoS scheduler policy	Basic Configurations
	Service management	Service Management Tasks
Slope QoS policies	Configure a slope QoS policy	Basic Configurations
	Service management	Service Management Tasks
Shared-Queue QoS policies	Configure a shared-queue QoS policy	Basic Configurations
Advanced QoS policies	Configure advanced QoS policies	Basic Configurations
	Service management	Service Management Tasks
QoS ATM traffic descriptor policies	Configure an ATM QoS traffic descriptor profile	Basic Configurations
	Service management	Service Management Tasks
Named pools	Configure a named pool QoS policy	Basic Configurations
HSMDA	Configure HSMDA	Basic HSMDA Configurations

Table 2 Configuration Process (Continued)

Area	Task	Section
High Scale QoS IOM	Configure HSQ IOM	High Scale QoS IOM
MC-MLPPP	Configure an MLPPP profile	Basic Configurations
Class Fair Hierarchal Policing (CHFP)	Configure CHFP	Class Fair Hierarchical Policing (CFHP)

2 QoS Policies

2.1 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 routers can classify, police, shape, and mark traffic.

In the Nokia service router service model, a service is provisioned on the provider-edge (PE) equipment. Service data is encapsulated, then sent in a service tunnel to the far-end Nokia 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 Nokia service routers (such as the 7950 XRS, 7750 SR, 7750 SR MG, and 7450 ESS) appears 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 described in more detail in [Forwarding Classes](#).

Router QoS policies 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 that 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.

2.2 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 how the queue output is scheduled into the switch fabric. The forwarding class of the packet, along with the 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 forwarding classes.

[Table 3](#) lists the default definitions for the forwarding classes. The forwarding class behavior, in terms of ingress marking interpretation and egress marking, can be changed by [Network QoS Policies](#). All forwarding class queues support the concept of in-profile, out-of-profile and, at egress only, inplus-profile and exceed-profile.

Table 3 Forwarding Classes

FC-ID	FC Name	FC Designation	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		

The forwarding classes can be classified into three class types:

- High-priority/Premium
- Assured
- Best-Effort

2.2.1 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](#)).

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 oversubscribed, a mechanism to engineer a path through the core network and to 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 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 using traffic engineered paths, as long as the core treats high-priority traffic with the correct PHB.

2.2.2 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 from 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. All DSCP and EXP mappings can be modified by the user.

2.2.3 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 by default as out-of-profile assured service packets.

2.3 Queue Parameters

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

2.3.1 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.

2.3.2 Unicast or Multipoint Queue

Unicast queues are used for all services and routing when the traffic is forwarded to a single destination.

Multipoint ingress queues are used by VPLS services for multicast, broadcast, and unknown traffic and by IES and VPRN services for multicast traffic when IGMP, MLD, or PIM is enabled on the service interface.

2.3.3 Queue Type

The type of a queue dictates how it will be scheduled relative to other queues at the hardware level. Being able to define the scheduling properties of a queue is important because a single queue allows support for multiple forwarding classes.

The queue type defines the relative priority of the queue and can be configured to **expedited** (higher priority) or **best-effort** (lower) priority. However, the instantaneous scheduling priority of a queue changes dynamically depending on its current scheduling rate compared to its operational CIR and FIR (see [Queue Scheduling](#)). Parental virtual schedulers can be defined for the queue using scheduler policies which enforce how the queue interacts for bandwidth with other queues associated with the same scheduler hierarchy (see [Scheduler Policies](#)).

The queue type of SAP ingress and egress queues, network queue policy queues, and ingress queue group template queues are defined at queue creation time. The queue type of egress queue group template queues and shared-queue policy queues can be modified after the queue has been created.

The default behavior for SAP ingress and egress queues, network queue policy queues, and shared queue policy queues is to automatically choose the expedited or best effort nature of the queue based on the forwarding classes mapped to it. This is achieved by configuring the queue type to **auto-expedited**. 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 best effort forwarding classes are mapped to the queue (be, af, l1, or l2), the queue will be treated as best effort by the hardware schedulers.

The default queue type for ingress queue group template and egress queue group template queues is **best-effort**.

2.3.4 Queue Scheduling

Packets are scheduled from queues by hardware schedulers based on the type of the queue (see [Queue Type](#)) and the current scheduling rate of the queue compared to its operational CIR and FIR. This applies to unicast queues at both ingress and egress, and multipoint queues at ingress, but not to HSMDB or HSQ IOM queues.

The queue type should be chosen based on the kind of traffic in the forwarding classes mapped to the queue.

The hardware scheduler services queues to forward packets from them in a strict priority order, as follows:

- Ingress queues, in priority order:
 1. Expedited queues where the queue's current scheduling rate is below its operational FIR
 2. Best effort queues where the queue's current scheduling rate is below its operational FIR
 3. Expedited queues where the queue's current scheduling rate is below its operational CIR
 4. Best effort queues where the queue's current scheduling rate is below its operational CIR
 5. Expedited queues where the queue's current scheduling rate is above both its operational FIR and CIR
 6. Best effort queues where the queue's current scheduling rate is above both its operational FIR and CIR
- Egress queues, in priority order:
 1. Expedited queues where the queue's current scheduling rate is below its operational FIR

2. Best effort queues where the queue's current scheduling rate is below its operational FIR
3. Expedited queues where the queue's current scheduling rate is below its operational CIR
4. Best effort queues where the queue's current scheduling rate is below its operational CIR
5. Expedited queues where the queue's current scheduling rate is above both its operational FIR and CIR
6. Best effort queues where the queue's current scheduling rate is above both its operational FIR and CIR



Note: Only network queues at egress support a non-zero FIR and can use priorities 1 and 2.

2.3.5 Peak Information Rate

The PIR defines the maximum rate at which packets are allowed to exit the queue. The PIR 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, FIR, and the queue type (see [Queue Type](#)) all combine to affect a queue's ability to transmit packets, as discussed in [Queue Scheduling](#).

The PIR is provisioned on ingress and egress queues within service ingress and egress QoS policies, network queue policies, ingress and egress queue group templates, and shared queue policies.

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 PIR value specified. The interpretation of the administrative PIR is discussed in [Adaptation Rule](#).

2.3.6 Committed Information Rate

The CIR for a queue performs two distinct functions:

1. Profile marking by service ingress queues — Service ingress queues (configured in SAP ingress QoS policies or ingress queue group templates) mark packets in-profile or out-of-profile based on the CIR. For each packet in a service ingress queue, the CIR is compared to 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. This operation can be overridden by configuring **cir-non-profiling** under the queue. This allows the queue scheduling priority to continue to be based on the below CIR or above CIR of the queue, but packets are not re-profiled depending on the state of the queue when they are scheduled (below CIR or above CIR). Instead, their profile state remains as out-of-profile, unless they are explicitly classified as in-profile or out-of-profile in which case they will remain in-profile or out-of-profile.
2. Scheduler queue priority — The scheduler serving a group of service ingress or egress queues prioritizes individual queues based on their current FIR, CIR, and PIR states. For more information about queue scheduling, see [Queue Scheduling](#).

All router queues support the concept of in-profile, out-of-profile, together with inplus-profile and exceed-profile at egress only. 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, the packets are dropped preferentially at congestion points in the core as follows:

- exceed-profile
- out-of-profile
- in-profile
- inplus-profile

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 if the hardware does not support the exact CIR specified. For more information about the interpretation of the administrative CIR, see [Adaptation Rule](#).

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 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 if this behavior is required. If the service queue is associated with a best-effort class, the CIR threshold is normally set to zero; however, this is flexible.

The CIR for a service queue is provisioned in ingress and egress service queues within service ingress QoS policies and service egress QoS policies respectively. CIRs for network queues are defined within network queue policies. CIRs for queue group instance queues are defined within ingress and egress queue group templates.

2.3.7 Fair Information Rate

A Fair Information Rate (FIR) can be configured on a queue to utilize two additional scheduling priorities (expedited queue within FIR and best-effort queue within FIR). The queues which operate below their FIR are always served before the queues operating at or above their FIR. For more information about queue scheduling, see [Queue Scheduling](#).

The FIR does not affect the queue packet profiling operation which is dependent only on the queue's CIR rate at the time that the packet is scheduled; the FIR is purely a scheduling mechanism.

When defining the FIR for a queue, the value specified is the administrative FIR for the queue. The router has a number of native rates in hardware that it uses to determine the operational FIR for the queue. The user has some control over how the administrative FIR is converted to an operational FIR if the hardware does not support the exact FIR specified. For more information about the interpretation of the administrative FIR, see [Adaptation Rule](#).

The FIR rate is supported in SAP ingress QoS policies (for ingress SAP and subscriber queues), ingress queue group templates (for access ingress queue group queues), and network queue policies (for ingress and egress network queues). The FIR can also be configured on system-created ingress shared queues. The FIR is configured in kb/s in a SAP ingress QoS policy and an ingress queue group template, and as a percent in a SAP ingress QoS policy, network queue policy, and ingress shared queue. The default FIR rate is zero (0) with an adaptation rule **closest**.

FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

The FIR rate is not used within virtual hierarchical scheduling (see [Virtual Hierarchical Scheduling](#)), but it is capped by the resulting operational PIR.

A FIR should not be configured in a SAP ingress QoS policy, ingress queue group template, or network queue policy associated with a LAG which spans FP4-based and FP2- or FP3-based hardware as the resulting operation could be different depending on which hardware type the traffic uses.

2.3.8 Adaptation Rule

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

For the FIR, 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 on which the queue is provisioned, the operational FIR, CIR, and PIR settings used by the queue will depend on the method that the hardware uses to implement and represent the mechanisms that enforce the FIR, CIR, and PIR rates.

As the hardware has a very granular set of rates, Nokia's recommended method to determine which hardware rate is used for a queue is to configure the queue rates with the associated adaptation rule and use the **show pools output** command to display the rate achieved.

To illustrate how the adaptation rule constraints (minimum, maximum, and closest) are evaluated in determining the operational CIR or PIR rates, assume there is a queue on an IOM3-XP where the administrative CIR and PIR rates are 401 Mb/s and 403 Mb/s, respectively.

The following output shows the operating CIR and PIR rates achieved for the different adaptation rule settings:

```
*A:PE# # queue using default adaptation-rule=closest
*A:PE# show qos sap-egress 10 detail | match expression "Queue-Id|CIR Rule"
Queue-Id          : 1          Queue-Type          : auto-expedite
PIR Rule          : closest    CIR Rule              : closest
*A:PE#
*A:PE# show pools 1/1/1 access-egress service 1 | match expression "PIR|CIR"
Admin PIR         : 403000      Oper PIR             : 403200
Admin CIR         : 401000      Oper CIR             : 401600
*A:PE#
*A:PE# configure qos sap-egress 10 queue 1 adaptation-rule pir max cir max
*A:PE#
*A:PE# show qos sap-egress 10 detail | match expression "Queue-Id|CIR Rule"
Queue-Id          : 1          Queue-Type          : auto-expedite
PIR Rule          : max        CIR Rule              : max
*A:PE#
*A:PE# show pools 1/1/1 access-egress service 1 | match expression "PIR|CIR"
Admin PIR         : 403000      Oper PIR             : 401600
Admin CIR         : 401000      Oper CIR             : 400000
*A:PE#
*A:PE# configure qos sap-egress 10 queue 1 adaptation-rule pir min cir min
*A:PE#
*A:PE# show qos sap-egress 10 detail | match expression "Queue-Id|CIR Rule"
Queue-Id          : 1          Queue-Type          : auto-expedite
PIR Rule          : min        CIR Rule              : min
*A:PE#
*A:PE# show pools 1/1/1 access-egress service 1 | match expression "PIR|CIR"
Admin PIR         : 403000      Oper PIR             : 403200
Admin CIR         : 401000      Oper CIR             : 401600
*A:PE#
```

2.3.9 Committed Burst Size

The CBS parameter specifies the size of buffer that can be drawn from the reserved buffer portion of the queue buffer pool. When the reserved buffers for a given queue have been used, the queue competes with other queues for additional buffer resources up to the MBS.

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

2.3.10 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 service queue is specified in bytes or kbytes.

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

2.3.11 Queue Drop Tails

The MBS determines the maximum queue depth after which no additional packets are accepted into the queue. Additional queue drop tails are available for the different packet profiles to allow preferential access to the queue's buffers which allows higher priority packets to be accepted into a queue when there is congestion for lower priority packets.


At ingress there is a low drop tail in addition to the MBS. High enqueueing priority packets (for ingress SAP priority mode queues) and in-profile packets (for ingress SAP profile mode queues, network, and shared queues) are allowed to fill up the queue up to the MBS, however, low enqueueing priority packets (for ingress SAP priority mode queues) and out-of-profile packets (for ingress SAP profile mode queues, network, and shared queues) can only fill the queue up to the queue's low drop tail setting.

At egress there are four drop tails in addition to the MBS, one for each profile state:

- an exceed drop tail for exceed-profile packets
- a low drop tail for out-of-profile packets
- a high drop tail for in-profile packets
- a highplus drop tail for highplus-profile packets

Each profile type can only fill the queue up to its corresponding drop tail. [Figure 1](#) shows the ingress and egress queue drop tails.

Figure 1 Ingress and Egress Queue Drop Tails



At both ingress and egress, the drop tails are configured as a percentage reduction from the MBS (specifying 10% places the drop tail at 90% of the MBS) and consequently all are limited by the queue's MBS.

The default percentage for the low drop tail for ingress SAP, queue group, and shared queues is a reduction from the MBS of:

- low = 10%

The default percentages for the drop tails for egress SAP, queue group, and network queues is a percentage reduction from the MBS of:

- exceed = 20%
- low = 10%
- high = 0%
- highplus = 0%

The exceed, high, and highplus drop tails are not configurable for network queues, however the exceed drop tail is set to a value of 10% in addition to low drop tail and capped by the MBS.

The four drop tails can be configured in any order within egress SAP and queue group queues, however it is logical to order them (from shortest to longest) as exceed, then low, then high, then highplus.

The low drop tail configuration can be overridden for ingress and egress SAP and queue group queues, and for network egress queues. It is also possible to override the low drop tail for subscriber queues within an SLA profile using the keyword **high-prio-only**.

When there is congestion the drop tail ordering gives preferential access to the queue's buffers. For example, if the drop tails on an egress SAP queue are configured as exceed = 20%, low = 10%, high = 5% and highplus = 0%, when the queue depth is below 80% all profile packet types are accepted into the queue. If the depth increases above 80%, then exceed profile packets will not be accepted and are therefore dropped, while the out-of-profile, in-profile, and inplus profile packets are still accepted into the queue (giving them preference over the exceed profile packets). If the queue depth goes beyond 90% the out-of-profile packets will also be dropped. Similarly, if the queue depth goes beyond 95% the in-profile packets will be dropped. It is only when the MBS has been reached that the inplus profile packets are dropped. This example assumes that the pool in which the queue exists is not congested.

2.3.12 WRED Per Queue

Egress SAP, subscriber, and network queues by default use drop tails within the queues and WRED slopes applied to the pools in which the queues reside in order to apply congestion control to the traffic in those queues. An alternative to this is to apply the WRED slopes directly to the egress queue using WRED per queue. See [Table 4](#).

WRED per queue is supported for SAP egress QoS policy queues (and therefore to egress SAP and subscriber queues) and for queues within an egress queue group template. There are two modes available:

- Native, which is supported on FP3- and higher-based hardware
- Pool-per-queue

Table 4 **WRED-Per-Queue Congestion Control Summary**


	Inplus-profile Traffic	In-profile Traffic	Out-of-profile Traffic	Exceed-profile Traffic	WRED-queue Mode	Slope Usage	Comments
FP3- and higher-based queues	Drop tail (MBS) used	Drop tail (MBS) used	Low slope used	Exceed slope used	Native	Exceed-low	If slope shutdown, MBS is used Highplus and high slope not used
Pool-per-queue	Highplus slope	High slope	Low slope	Exceed slope	Pool-per-queue	Default	If slope shutdown, MBS is used
Pool/megapool/named-pool	Highplus slope	High slope	Low slope	Exceed slope	N/A	N/A	If slope shutdown, total pool size is used

2.3.12.1 Native Queue Mode

When an egress queue is configured for native mode, it will use the native WRED capabilities of the forwarding plane queue. This is supported on FP3- and higher-based hardware.

Congestion control within the queue will use the low and exceed slopes from the applied slope policy together with the MBS drop tail. The queue continues to take buffers from its associated egress access or network buffer pool, on which WRED can also be enabled. This is shown in [Figure 2](#).

Figure 2 **WRED Queue: Native Mode**



To configure a native WRED queue, the **wred-queue** command is used under **queue** in a SAP egress QoS policy or egress queue group template with the mode set to **native** as follows:

```
qos
queue-group-templates
  egress
    queue-group <queue-group-name> create
    queue <queue-id> create
      wred-queue [policy <slope-policy-name>] mode native
      slope-usage exceed-low
    exit
  exit
exit
sap-egress <policy-id> create
  queue <queue-id> [<queue-type>] [create]
    wred-queue [policy <slope-policy-name>] mode native
    slope-usage exceed-low
  exit
exit
```

Congestion control is provided by both the slope policy applied to the queue and the MBS drop tail.

The **slope-usage** defines the mapping of the traffic profile to the WRED slope and only **exceed-low** is allowed with a native mode queue. The slope mapping is shown in the following example and requires the low and exceed slopes to be no shutdown in the applied slope policy (otherwise traffic will use the MBS drop tail or a pool slope):

- Out-of-profile maps to the low slope.
- Exceed-profile maps to the exceed slope.

The instantaneous queue depth is used against the slopes when native mode is configured; consequently, the time-average-factor within the slope policy is ignored.

The inplus and in-profile traffic uses the MBS drop tail for congestion control (the highplus or high slopes are not used with a native mode queue).

If a queue is configured to use native mode WRED per queue on hardware earlier than FP3, the queue operates as a regular queue.

For example, the following SAP egress QoS policy is applied to SAP on an FP3- or higher-based hardware:

```
sap-egress 10 create
  queue 1 create
```

```
wred-queue policy "slope1" mode native
    slope-usage exceed-low
exit
exit
```

The details of both the pool and queue can then be shown using the following command:

```
A:PE# show pools 1/1/1 access-egress service 1
=====
Pool Information
=====
Port                : 1/1/1
Application         : Acc-Egr           Pool Name           : default
CLI Config. Resv CBS : 30%(default)
Resv CBS Step       : 0%                Resv CBS Max         : 0%
Amber Alarm Threshold: 0%                Red Alarm Threshold  : 0%
-----
Queue-Groups
-----
Queue-Group:Instance : policer-output-queues:1
-----
Utilization          State      Start-Avg    Max-Avg      Max-Prob
-----
HiPlus-Slope         Down          85%          100%         80%
High-Slope           Down          70%          90%          80%
Low-Slope            Down          50%          75%          80%
Exceed-Slope         Down          30%          55%          80%
Time Avg Factor      : 7
Pool Total           : 192000 KB
Pool Shared          : 134400 KB          Pool Resv            : 57600 KB
-----
Current Resv CBS     Provisioned   Rising        Falling        Alarm
%age                 all Queues   Alarm Thd     Alarm Thd      Color
-----
30%                  0 KB        NA            NA            Green
Pool Total In Use    : 0 KB
Pool Shared In Use   : 0 KB          Pool Resv In Use     : 0 KB
WA Shared In Use     : 0 KB
HiPlus-Slope Drop Pr*: 0          Hi-Slope Drop Prob   : 0
Lo-Slope Drop Prob   : 0          Excd-Slope Drop Prob  : 0
=====
Queue : 1->1/1/1->1
=====
FC Map              : be 12 af 11 h2 ef h1 nc
Tap                 : not-applicable    FP                 : not-applicable
Admin PIR           : 10000000         Oper PIR           : Max
Admin CIR           : 0                Oper CIR           : 0
Admin MBS           : 12000 KB         Oper MBS           : 12000 KB
High-Plus Drop Tail : not-applicable    High Drop Tail     : not-applicable
Low Drop Tail       : not-applicable    Exceed Drop Tail   : not-applicable
CBS                 : 0 KB             Depth              : 0
Slope               : slope1
WRED Mode           : native           Slope Usage        : exceed-low
-----
HighPlus Slope Information
```

```

-----
State           : disabled
Start Average   : 0 KB
Max Probability  : 0 %
Max Average     : 0 KB
Curr Probability : 0 %
-----

High Slope Information
-----
State           : disabled
Start Average   : 0 KB
Max Probability  : 0 %
Max Average     : 0 KB
Curr Probability : 0 %
-----

Low Slope Information
-----
State           : enabled
Start Average   : 6960 KB
Max Probability  : 80 %
Max Average     : 8880 KB
Curr Probability : 0 %
-----


Exceed Slope Information
-----
State           : enabled
Start Average   : 4680 KB
Max Probability  : 80 %
Max Average     : 6600 KB
Curr Probability : 0 %
-----
=====
* indicates that the corresponding row element may have been truncated.
A:PE#

```

2.3.12.2 Pool Per Queue Mode

When pool per queue mode is used, the queue resides in its own pool that is located in the forwarding plane egress megapool. The size of the pool is the same as the size of the queue (based on the MBS); consequently, the WRED slopes operating on the pool's buffer utilization are reacting to the congestion depth of the queue. The size of the reserved CBS portion of the buffer pool is dictated by the queue's CBS parameter. This is shown in [Figure 3](#).

Figure 3 WRED Queue: Pool-per-Queue Mode



0947

The queue pools take buffers from the WRED egress megapool that must be enabled per FP; if this megapool is not enabled, the queue operates as a regular queue. By default, only the ingress normal and egress normal megapools exist on an FP. The egress WRED megapool is configured using the following commands:

```
configure
card <slot-number>
fp <fp-number>
egress
wred-queue-control
[no] buffer-allocation min <percentage> max
<percentage>
[no] resv-cbs min <percentage> max <percentage>
[no] shutdown
[no] slope-policy <slope-policy-name>
```

The buffer allocation determines how much of the egress normal megapool is allocated for the egress WRED megapool, with the **resv-cbs** defining the amount of reserved buffer space in the egress WRED megapool. In both cases, the **min** and **max** values must be equal. The **slope-policy** defines the WRED slope parameters and the time average factor used on the megapool to handle congestion as the megapool buffers are used. The **no shutdown** command enables the megapool. The megapools on card 1 fp 1 can be shown as follows:

```
*A:PE# show megapools 5 fp 1
=====
MegaPool Summary
=====
```

Type	Card-FP	App.	Pool Name	Actual Admin	ResvCBS ResvCBS	PoolSize
fp	5/1	Egress	normal	n/a	n/a	282624
fp	5/1	Ingress	normal	n/a	n/a	374784

fp	5/1	Egress wred	n/a	
			24576	93696
			25%	

=====

*A:PE#

To configure a pool per queue, the **wred-queue** command is used under **queue** in a SAP egress QoS policy or egress queue group template with the mode set to **pool-per-queue** as follows:

```

qos
queue-group-templates
  egress
    queue-group <queue-group-name> create
    queue <queue-id> create
      wred-queue [policy <slope-policy-name>] mode pool-
per-queue
      slope-usage default
    exit
  exit
exit
sap-egress <policy-id> create
  queue <queue-id> [<queue-type>] [create]
    wred-queue [policy <policy-name>] mode pool-per-queue
    slope-usage default
  exit
exit

```

Congestion control is provided by the slope policy applied to the queue, with the slopes to be used having been **no shutdown** (otherwise traffic will use the MBS drop tail or a megapool slope). The **slope-usage** defines the mapping of the traffic profile to the WRED slope and only **default** is allowed with **pool-per-queue** that gives the following mapping:

- Inplus-profile maps to the highplus slope.
- In-profile maps to the high slope.
- Out-of-profile maps to the low slope.
- Exceed-profile maps to the exceed slope.

For example, the following SAP egress QoS policy is applied to SAP on an FP with the egress WRED megapool enabled:

```

sap-egress 10 create
  queue 1 create
    wred-queue policy "slope1" mode pool-per-queue
    slope-usage default
  exit

```


exit

The details of both the megapool and queue pool usages can then be shown using the **show megapools** command (detailed version):

```
*A:PE# show megapools 5 fp 1 wred service-id 1
=====
MegaPool Information
=====
Slot                : 5                FP                : 1
Application         : Egress           Pool Name          : wred
Resv CBS            : 25%
-----
Utilization          State             Start-Avg           Max-Avg             Max-Prob
-----
HiPlus-Slope         Down                85%                 100%                80%
High-Slope           Down                70%                 90%                 80%
Low-Slope             Down                50%                 75%                 80%
Exceed-Slope         Down                30%                 55%                 80%
Time Avg Factor      : 7
Pool Total           : 93696 KB
Pool Shared          : 69120 KB          Pool Resv            : 24576 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
HiPlus-Slope Drop Pr*: 0             Hi-Slope Drop Prob  : 0
Lo-Slope Drop Prob   : 0             Excd-Slope Drop Prob : 0
=====
Queue : 1->5/1/1:1->1
=====
FC Map               : be l2 af l1 h2 ef h1 nc
Tap                  : not-applicable
Admin PIR             : 10000000          Oper PIR             : Max
Admin CIR             : 0                Oper CIR             : 0
Admin MBS             : 12288 KB          Oper MBS             : 12288 KB
High-Plus Drop Tail  : not-applicable   High Drop Tail       : not-applicable
Low Drop Tail         : not-applicable   Exceed Drop Tail     : not-applicable
CBS                   : 0 KB             Depth                : 0
Slope                 : slop1
WRED Mode             : pool-per-queue   Slope Usage          : default
Time Avg Factor       : 7
-----
HighPlus Slope Information
-----
State                : enabled
Start Average        : 10752 KB          Max Average          : 12288 KB
Max Probability      : 80 %              Curr Probability     : 0 %
-----
High Slope Information
-----
State                : enabled
Start Average        : 9408 KB           Max Average          : 10944 KB
Max Probability      : 80 %              Curr Probability     : 0 %
-----
Low Slope Information
-----
```

```

State           : enabled
Start Average   : 6144 KB           Max Average      : 9216 KB
Max Probability : 80 %              Curr Probability : 0 %
-----
Exceed Slope Information
-----
State           : enabled
Start Average   : 3648 KB           Max Average      : 6720 KB
Max Probability : 80 %              Curr Probability : 0 %
=====
* indicates that the corresponding row element may have been truncated.
*A:PE#

```

Each WRED **pool-per-queue** uses a WRED pool resource on the FP. The resource usage is displayed in the **tools dump resource-usage card [slot-num] fp [fp-number]** command output under the Dynamic Q2 WRED Pools header.

2.3.13 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 non-zero low drop tail setting, 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 ingress queue for the packet is above the low drop tail setting, 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. After 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.

At access ingress, the priority of a packet has no effect 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 and FIR (where supported) does affect the scheduling priority between queues going to the same destination (either the switch fabric tap or egress port). For information about the strict operating priority for queues, see [Queue Scheduling](#).

For access ingress, the CIR controls both dynamic scheduling priority and marking threshold (unless **cir-non-profiling** is configured). 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 determines the dynamic scheduling priority, but does not affect the packet's ingress determined profile.

2.3.14 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

2.3.15 Color Aware Profiling

The normal handling of SAP ingress access packets applies an in-profile or out-of-profile state (associated with the colors green and yellow, respectively) to each packet relative to the dynamic rate of the queue while the packet is forwarded toward 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. When 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 toward 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 subclass associations for SAP ingress classification purposes. Each subclass retains the chassis-wide behavior defined to the parent class while providing expanded ingress QoS classification actions. Subclasses 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.

2.4 QoS Policies Overview

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; 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 resource-usage card fp** output displays the number of policies applied on an FP (the default SAP ingress policy is always applied once for internal use). The **tools dump resource-usage system** output displays the usage of the policies at a system level. The **show qos sap-ingress** and **show qos sap-egress** commands can be used to show the number of policies 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

The router supports thousands of queues. The exact numbers depend on the hardware being deployed.

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 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 criteria, 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 unicast forwarding class queues in the policy; one for each forwarding class. A service ingress QoS policy also defines up to three queues per forwarding class to be used for multipoint traffic for multipoint services.

In the case of VPLS, four types of forwarding are supported (that 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 manner within the service.

Service egress QoS policies are applied to SAPs and map forwarding classes to service egress queues for a service. Up to eight queues per service can be defined for the eight 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 XMAAs or 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 5](#).

Table 5 QoS Policy Types and Descriptions


Policy Type	Applied at...	Description	See
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/DE bits, DSCP, TOS precedence), IP criteria, or MAC criteria	Service Ingress QoS Policies
Service Egress	SAP egress	Defines up to eight forwarding class queues and queue parameters for traffic classification Maps one or more forwarding classes to the queues	Service Egress QoS Policies
Network	Router interface	Used for classification/marketing of IP and MPLS packets At ingress, defines DSCP, dot1p MPLS LSP-EXP, and IP criteria classification to FC mapping At ingress, defines FC to policer/queue-group queue mapping At egress, defines DSCP or precedence FC mapping At egress, defines FC to policer/queue-group queue mapping At egress, defines DSCP, MPLS LSP-EXP, and dot1p/DE marking	Network QoS Policies
Network Queue	Network ingress MDA and egress port	Defines forwarding class mappings to network queues and queue characteristics for the queues	Network Queue QoS Policies
Slope	Ports	Enables or disables the WRED slope parameters within an egress queue, an ingress pool, or an egress megapool	Slope Policy Parameters
Scheduler	Customer multiservice site Service SAP	Defines the hierarchy and parameters for each scheduler Defined in the context of a tier that is used to place the scheduler within the hierarchy Three tiers of virtual schedulers are supported	Scheduler Policies

Table 5 QoS Policy Types and Descriptions (Continued)

Policy Type	Applied at...	Description	See
Shared Queue	SAP ingress	Shared queues can be implemented to mitigate the queue consumption on an MDA	Shared-Queue QoS Policies
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	ATM Traffic Descriptor Profiles
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	ATM Traffic Descriptor Profiles

2.4.1 Service versus Network QoS

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, as shown in [Figure 4](#).

Figure 4 Service vs Network Traffic Types

The router uses QoS policies applied to a SAP for a service or to a network MDA/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

2.4.2 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 a new QoS policy is created, 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, description, and at least one queue. The queue is associated with a forwarding class.

Service QoS policies can be applied to the following service types:

- Epipe — Both ingress and egress policies are supported on an Epipe 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.

Network QoS policies can be applied to the following entities:

- Network interfaces
- Apipe/Cpipe/Epipe/Ipipe/Fpipe and VPLS spoke-SDPs
- VPLS spoke- and mesh-SDPs
- IES and VPRN interface spoke-SDPs
- VPRN network ingress
- VXLAN network ingress

Network QoS policies can be applied to:

- Ingress MDAs
- Egress ports

Default QoS policies map 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.

2.4.3 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, Dot1p, and IP criteria mappings to forwarding classes
 - Defines LSP EXP value mappings to forwarding classes
- Egress
 - Defines DSCP and IP precedence mappings to forwarding classes
 - Defines forwarding class to DSCP value markings
 - Defines forwarding class to Dot1p/DE value markings
 - Defines forwarding class to LSP EXP value markings
 - Enables/disables remarking of QoS
 - Defines FC to policer/queue-group queue mapping

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 Dot1p 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
- Ingress FC fp-redirect-group policer mapping
- Egress FC port-redirect-group queue/policer mapping

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 that do not have another network QoS policy explicitly assigned. [Table 6](#) describes the default network QoS policy egress marking.

Table 6 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 7](#) and [Table 8](#) list the default mapping of DSCP name and LSP EXP values to forwarding class and profile state for the default network QoS policy.

Table 7 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		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
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

2.4.4 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 or 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 the same number of queues as the default network queue policy. The multicast queues are only used at ingress.

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

- CBS as a percentage of the buffer pool
- MBS as a percentage of the buffer pool
- Low drop tail as a percentage reduction from MBS
- PIR as a percentage of the FP ingress capacity or egress port bandwidth
- CIR as a percentage of the FP ingress capacity or egress port bandwidth
- FIR as a percentage of the FP ingress capacity or egress port bandwidth

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

The system default network queue policy is named **default** and cannot be edited or deleted. For information about the default network queue policy, see [Network Queue QoS Policies](#).

2.4.5 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](#).
- At least one multipoint forwarding class queue

Optional service ingress QoS policy elements include:

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

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

There can 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 subclasses. A portion of the subclass definition directly ties the subclass to a parent, chassis-wide forwarding class. The subclass is only used as a SAP ingress QoS classification tool; the subclass association is lost when ingress QoS processing is finished.

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 8](#) lists the classification hierarchy based on rule type.

Table 8 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 must be explicitly defined. Each packet can only match a single prec rule.

Table 8 Forwarding Class and Enqueuing Priority Classification Hierarchy Based on Rule Type (Continued)

#	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, 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, then processing stops. A packet can only match a single MAC criteria entry.

2.4.5.1 FC Mapping Based on EXP Bits at VLL/VPLS SAP

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 Ethernet Layer 2 SAPs. The traffic will be encapsulated in a VLL/VPLS SDP. These small ISP routers are typically PE routers. To provide appropriate QoS, the routers support a new classification option that based on received MPLS EXP bits. [Figure 5](#) shows a sample configuration.

Figure 5 Example Configuration — Carrier Application

OSSG257

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

[Table 9](#) lists forwarding class behavior by rule type.

Table 9 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, 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, then processing stops. A packet can only match a single MAC criteria entry.

The enqueueing priority is specified as part of the classification rule and is set to “high” or “low”. The enqueueing priority relates to the forwarding class queue’s low drop tail where only packets with a high enqueueing priority are accepted into the queue when the queue’s depth reaches the defined threshold. See [Queue Drop Tails](#).

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 that specifies:

- The forwarding class of packets that match the entry
- The enqueueing 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.

The supported service ingress QoS policy IP match criteria are:

- Destination IP address/prefix
- Destination port/range
- IP fragment
- Protocol type (TCP, UDP, and so on)
- Source port/range
- Source IP address/prefix
- DSCP value

The supported service ingress QoS policy MAC match criteria are:

- 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 10](#).

Table 10 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 2-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 11](#) lists the criteria that can be matched for the various MAC frame types.

Table 11 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 ¹	Yes	Yes
ethernet-II	Yes	Yes	Yes	Yes	No	No	No

Note:

1. 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 that do not explicitly have another service ingress policy assigned. The characteristics of the default policy are listed in [Table 12](#).

Table 12 **Default Service Ingress Policy ID 1 Definition**

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


2.4.5.2 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 subclass) and by default the same forwarding class is used for this traffic at the access or network egress. The ingress forwarding class or subclass 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 subclasses, 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 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 6 Egress Forwarding Class Override

al_0187

Figure 6 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.

2.4.6 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 eight separate queues (unicast)
- Dot1p/DE, DSCP, and IP precedence 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 classes 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 either at ingress or egress. 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 on a network interface, the forwarding class is marked in the tunnel transport encapsulation. In each case, the packet can be reclassified into a different forwarding class at service egress.

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 that do not have another service egress policy explicitly assigned. The characteristics of the default policy are listed in [Table 13](#).

Table 13 **Default Service Egress Policy ID 1 Definition**

Characteristic	Item	Definition
Queues	Queue 1	One queue defined for all traffic classes: • CIR = 0 • PIR = max (line rate) • MBS and CBS = default
Flows	Default Action	One flow defined for all traffic classes: • All traffic mapped to queue 1 with no marking

2.4.7 Named Pool Policies

The named buffer pool feature for the 7450 ESS and 7750 SR allows for the creation of named buffer pools at the MDA and port level. Named pools allow for a customized buffer allocation mode for ingress and egress queues that goes beyond the default pool behavior.

Named pools are defined within a named pool policy. The policy contains a q1-pools context that is used to define port allocation weights and named pools for buffer pools on Q1-based IOMs (all IOMs that are supported). The policy may be applied at either the port or MDA level at which time the pools defined within the policy are created on the port or MDA. When the policy is applied at the MDA level, MDA named pools are created.

MDA named pools will typically be used when either a pool cannot be created per port or when the buffering needs of queues mapped to the pool are not affected by sharing the pool with queues from other ports. MDA named pools allow buffers to be efficiently shared between queues on different ports mapped to the same pool. However, with MDA named pools, very active queues on one port could deplete buffers in the pool, with queues on other ports experiencing buffer starvation.

Port named pools are created when the policy is applied at the port level and allow for a more precise application of the buffer space allocated for a physical port. MDA pool names do not need to be unique. If name overlaps exist, the port pool will be used. The same pool name may be created on multiple ports on the same MDA.

The named pool policy is applied at the MDA ingress and egress level and at the ingress and egress port level. Each MDA within the system is associated with a forwarding plane traffic manager that has support for a maximum of 57 buffer pools. The following circumstances affect the number of named pools that can be created per MDA (these circumstances may be different between ingress and egress for the MDA):

- The forwarding plane can be associated with multiple MDAs (each MDA has its own named pools).
- A single system-level pool for system-created queues is allocated.
- There must be default pools for queues that are not explicitly mapped or are incorrectly mapped to a named pool.
- Default pools for most IOM types (separate for ingress and egress).
 - Access pool
 - Network pool
- The number of named per-port pools is dependent on the number of ports the MDA supports, which is variable per MDA type.

- Per-port named pools cannot be used by ingress network queues, but pools defined in a named pool policy defined on an ingress all network port are still created:
 - Ingress network queues use the default network pool or MDA named pools.
 - Ingress port buffer space allocated to network mode ports is included in the buffers made available to ingress MDA named pools.
 - Ingress port buffer space on channelized ports associated with network bandwidth is included in the buffers made available to ingress MDA named pools.
 - Ingress port named pools are only allocated buffers when the port is associated with some access mode bandwidth.
- Per-port named pools on ports aggregated into a LAG are still created per physical port.
- Default, named MDA, and named per-port pools are allocated regardless of queue provisioning activity associated with the pool.

If the named pool policy is applied to an MDA or port that cannot create every pool defined in the policy, the policy application attempt will fail. Any preexisting named pool policy on the MDA or port will not be affected by the failed named pool policy association attempt.

When buffer pools are being created or deleted, individual queues may need to be moved to or from the default pools. When a queue is being moved, the traffic destined to the queue is first moved temporarily to a fail-over queue. Then the queue is allowed to drain. When the queue is drained, the statistics for the queue are copied. The queue is then returned to the free queue list. A new queue is created, associated with the appropriate buffer pool, the saved statistics are loaded to the queue, and the traffic is moved from the fail-over queue to the new queue.

While the traffic is being moved between the old queue to the fail-over queue, then to the new queue, some out of order forwarding may be experienced. Also, any traffic forwarded through the fail-over queue will not be accounted for in billing or accounting statistics. A similar action is performed for queues that have the associated pool name added, changed, or removed. This only applies to where fail-over queues are supported.

The first step in allowing named pools to be created for an MDA is to enable named pool mode at the IOM level (**config card slot-number named-pool-mode**). Named pool mode may be enabled and disabled at any time. When MDAs are currently provisioned on the IOM, the IOM is reset to allow all existing pools to be deleted and the new default, named MDA, and named port pools to be created and sized. If MDAs are not currently provisioned (as when the system is booting up), the IOM is not reset.

When named pool mode is enabled, the system changes the way that default pools are created. The system no longer creates default pools per port; instead, a set of per forwarding plane level pools are created that are used by all queues that are not explicitly mapped to a named pool.

After the IOM has been placed into named pool mode, a named pool policy must be associated with the ingress and egress contexts of the MDA or individual ports on the MDA for named pools to be created. There are no named pools that exist by default.

Each time the default pool reserve, aggregate MDA pool limit, or individual pool sizes is changed, buffer pool allocation must be reevaluated.

Pools may be deleted from the named pool policy at any time. Queues associated with removed or non-existent pools are mapped to one of the default pools based on whether the queue is access or ingress. The queue is flagged as pool-orphaned until either the pool exists, or the pool name association is changed on the pool.

An ingress or egress port-managed buffer space is derived from the port's active bandwidth. Based on this bandwidth value compared to the bandwidth value of the other ports, the available buffer space is given to each port to manage. It may be required to artificially increase or decrease this bandwidth value to compensate for how many buffers are needed on each port.

If one port has very few queues associated with it and another has many queues associated, the commands in the port's "modify-buffer-allocation-rate" CLI context may be used to move one port's bandwidth up, and another port's bandwidth down. As provisioning levels change between ports, the rate modification commands may be used to adapt the buffer allocations per port.

Buffer allocation rate modification is supported for both standard and named pool mode buffer allocation methods.

The system allocates buffers based on the following criteria:

- the "named-pool-mode" setting on the IOM
- the amount of path bandwidth on channelized ports
- the existence of queues provisioned on the port or channel
- the current speed of each port
- the "ing-percentage-of-rate" and "egr-percentage-of-rate" command setting of each port
- the port-allocation-weights setting for default, MDA, and port
- the port's division between network and access bandwidth
- each individual named pool's network-allocation-weight and access-allocation-weight

2.4.8 Slope Policies

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

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 XMA and MDA levels. Each physical port has two pool objects associated:

- Access ingress/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.

Slope policies can also be applied when using WRED per queue; see [WRED Per Queue](#).

2.4.8.1 RED Slopes

Each buffer pool supports a high-priority 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. In addition, egress access, network pools, and megapools support a highplus slope that manages access to the shared portion of the buffer pool for inplus-profile packets and an exceed slope that manages access to the shared portion of the buffer pool for exceed-profile packets.

For access and network buffer pools, the percentage of the buffers that are to be reserved for CBS buffers is configurable. 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.

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.


Packets that were either classified as high priority or considered in-profile are 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. At egress, the highplus slope should be configured with RED parameters that prioritize the inplus priority above the other priority traffic and the exceed-slope should be configured with RED parameters that prioritize the high-priority and low-priority traffic above the exceed-profile traffic.

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 (SBU) and Shared Buffer Average Utilization (SBAU).
2. At initialization, the SBU is zero and the SBAU is zero.
3. When each packet is received, the current SBAU 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 SBU size.
8. If the packet is queued, a new SBAU is calculated using the time-average-factor (TAF) for the buffer pool. The TAF describes the weighting between the previous SBAU result and the new SBU in determining the new SBAU. (See [Tuning the SBAU Calculation](#).)

9. The new SBAU is used as the SBAU the next time that 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 SBU 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 SBU.

Figure 7 RED Slope Characteristics



A RED slope 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%. 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 (see [Figure 7](#)):

1. Section A is (0, 0) to (start-avg, 0). This is the part of the slope where the packet discard value is always 0, 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 0 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 1. 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.

2.4.8.2 Tuning the SBAU Calculation

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

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

sw0533

where:

SBAUn = SBAU for event n

SBAUn - 1 = SBAU for event (n - 1)

SBU = Instantaneous SBU

TAF = Time average factor

[Table 14](#) describes the effect that the allowed values of TAF have on the relative weighting of the instantaneous SBU and the previous SBAU (SBAUn-1), when calculating the current SBAU (SBAUn).

Table 14 TAF Impact on Shared Buffer Average Utilization Calculation

TAF	2^{TAF}	Equates To	Instantaneous SBU Portion	SBAU 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)
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)	1023/1024 (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.

2.4.8.3 Slope Policy Parameters

The elements required to define a slope policy are:

- a unique policy ID
- the highplus, high, low, and exceed 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 port 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 that do not have another slope policy explicitly assigned.

[Table 15](#) lists the default values for the default slope policy.

Table 15 Default Slope Policy Definition

Parameter	Description	Setting
Policy ID	Slope policy ID	1 (Policy ID 1 reserved for default slope policy)
Highplus (RED) slope	Administrative state	Shutdown
	start-avg	85% utilization
	max-avg	100% utilization
	max-prob	80% probability
High (RED) slope	Administrative state	Shutdown
	start-avg	70% utilization
	max-avg	90% utilization
	max-prob	80% probability

Table 15 Default Slope Policy Definition (Continued)

Parameter	Description	Setting
Low (RED) slope	Administrative state	Shutdown
	start-avg	50% utilization
	max-avg	75% utilization
	max-prob	80% probability
Exceed (RED) slope	Administrative state	Shutdown
	start-avg	30% utilization
	max-avg	55% utilization
	max-prob	80% probability
TAF	Time average factor	7

2.4.9 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 created when the policy is applied to an object, such as a SAP or interface.

Scheduler objects define bandwidth controls that limit each child (other schedulers and policers or 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 policers, queues, or other schedulers defined as child associations. The scheduler can also be a child which takes bandwidth from a scheduler in a higher tier.


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. When multiple schedulers, policers, and 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 any time and is immediately reflected on the schedulers 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 policers or queues (where the specified scheduler name 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 8](#) shows how child queues, policers, 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 8 Virtual Scheduler Internal Bandwidth Allocation



al_0188

2.4.9.1 Virtual Hierarchical Scheduling

Virtual hierarchical scheduling is a method that defines a bounded operation for a group of queues and, at egress only, policers. One or more policers and queues are mapped to a 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 SLA to be enforced.

The virtual hierarchical scheduling bandwidth distribution mechanism complements the packet scheduling provided by the queue scheduling (see [Queue Scheduling](#)).

Scheduler policies in the routers determine how a policer or queue interacts with bandwidth with other children associated with the same scheduler hierarchy. Ingress queues and egress queues and policers can operate within the context of a scheduler. Multiple policers and queues can 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

2.4.9.2 Hierarchical Scheduler Policies

Hierarchical scheduler policies are an alternate way of scheduling that can be used on service ingress queues and service egress queues and policers. Hierarchical scheduler policies allow the creation of a hierarchy of schedulers where policers, 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 HQoS on the SR OS.

2.4.9.2.1 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 (see [Figure 9](#)).

Tier 1 schedulers (also called root schedulers) can also have a parent scheduler. A scheduler can enforce a maximum rate of operation for all child policers, queues, and associated schedulers.

Figure 9 Hierarchical Scheduler and Queue Association



2.4.9.3 Tiers

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 Mb/s, PIR = 10 Mb/s
- Silver: CIR = 20 Mb/s, PIR = 40 Mb/s

- Bronze: CIR = 0 Mb/s, PIR = 100 Mb/s

In the router, the CIR is used for profiling 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 Mb/s (10 Mb/s + 40 Mb/s + 100 Mb/s) 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 that limits the overall rate for all queues to 100 Mb/s. 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 Mb/s.

2.4.9.4 Scheduler Policies Applied to Applications

A scheduler policy can be applied either on a SAP (see [Figure 10](#)) or on a multiservice customer site (a group of SAPs with common origination/termination point) (see [Figure 11](#)). 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 policers and queues created on that SAP can use the schedulers created by the policy association. When the object is a multiservice customer site, the schedulers are available to any SAPs associated with the site (also see [Scheduler Policies Applied to SAPs](#)).

Figure 10 Scheduler Policy on SAP and Scheduler Hierarchy Creation





Figure 11 Scheduler Policy on Customer Site and Scheduler Hierarchy Creation



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

2.4.9.5 Scheduler Policies Applied to SAPs

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

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

2.4.9.6 Scheduler Policies Applied to Customer SLAs

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 accounted and billed as an aggregate, based on a single SLA.

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

When 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 preprovisioned (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. Multiservice customer sites are configured only to create a virtual scheduler hierarchy and make it available to queues on multiple SAPs.

2.4.9.7 Scheduler Policies Applied to Multiservice Sites

Only existing scheduler policies can be applied to create the ingress or egress schedulers used by SAP policers or queues associated with a customer's multiservice 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. When a multiservice customer site is created, SAPs owned by the customer must be explicitly included in the site. The SAP must be owned by the customer within which the site was created and the site assignment parameter must include the physical locale of the SAP.

2.4.10 ATM Traffic Descriptor Profiles

Traffic descriptor 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.

ATM traffic descriptor profiles are not supported on the 7950 XRS.

2.4.11 Configuration Notes

The following information describes QoS implementation restrictions:

- Creating additional QoS policies is optional.
- Default policies are created for service ingress, service egress, network, network queue, and 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.

3 Match List for QoS Policies

Match lists provide a mechanism to simplify the configuration of IP and IPv6 criteria matching statements within QoS policies. Instead of defining multiple match statements in an **ip-criteria** or **ipv6-criteria** statement, an operator can group the same types of matching criteria into a single match list and use that list as a match criterion value, thereby requiring only a single policy entry per each unique action. The same match list can be used in one or more QoS policies.

The match lists further simplify management and deployment of the policy changes. A change in a **match-list** content is automatically propagated across all policies employing that list in their match criteria, therefore, only a single configuration change is required to trigger policy changes when a list is used by entries in one or more QoS policies.

The hardware resource usage does not change whether QoS match lists are used or whether the operator creates multiple entries (one for each element in the list). However, consideration must be given to how the lists are used to ensure only needed match permutations are created in a QoS policy entry (especially when other match criteria that are also lists or ranges are specified in the same entry). The system verifies whether a new list element, for example, an IP address prefix, can be added to a specific list, or a list can be used by a new QoS policy, by checking whether the resources exist in hardware to implement the required changes for all QoS policies that reference the updated list. If sufficient resources do not exist, the addition of a new element to the list or use of the list by another policy will fail.

QoS match lists are created within **config>qos>match-list**. The following types of match lists are supported:

- IPv4 prefix lists – Applicable to **src-ip** and **dst-ip** matching in SAP ingress and SAP egress QoS policies used by both SAPs and subscribers, and in the ingress section of a network QoS policy.
- IPv6 prefix lists – Applicable to **src-ip** and **dst-ip** matching in the ingress section of a network QoS policy.

The following restrictions apply to the use of prefix lists in network QoS policies:

- A single IP prefix list (IPv4/IPv6) cannot be used by network QoS policy entries more than 128 times.
- A single entry in a network QoS policy can only refer to either a source or destination prefix list. It is not permitted to refer simultaneously to both a source and a destination prefix IPv4/IPv6 list.
- A given prefix list can be attached to a SAP or network QoS policy, but not both simultaneously.

The following shows a created IPv4 prefix list which is configured within a SAP ingress QoS policy to rate limit the traffic from those prefixes.

```
configure
#-----
echo "QoS Policy Configuration"
#-----
  qos
    match-list
      ip-prefix-list "ip-prefix-list-1" create
      description "IPv4 prefix list"
      prefix 10.0.0.0/8
      prefix 192.168.0.0/16
    exit
  exit
exit
#-----
echo "QoS Policy Configuration"
#-----
  qos
    sap-egress 10 create
    queue 1 create
    exit
    queue 2 create
    exit
    fc af create
    queue 2
    exit
    ip-criteria
      entry 10 create
      match
        dst-ip ip-prefix-list "ip-prefix-list-1"
      exit
      action fc "af"
    exit
  exit
exit
exit
```

The IPv4 prefix list can be shown as follows:

```
*A:PE# show qos match-list ip-prefix-list "ip-prefix-list-1"

=====
QoS Match IP Prefix List
=====
Prefix Name       : ip-prefix-list-1
Description       : IPv4 prefix list
-----
IP Prefixes
-----
10.0.0.0/8
192.168.0.0/16
-----
No. of Prefixes : 2
-----
=====
```

*A:PE#

3.1 Match List for QoS Policies Commands

3.1.1 Command Hierarchies

- [Configuration Commands](#)
- [Show Commands](#)

3.1.1.1 Configuration Commands

```
config
  — qos
    — match-list
      — ip-prefix-list ip-prefix-list-name [create]
      — no ip-prefix-list ip-prefix-list-name
        — description description-string
        — no description
        — [no] prefix ip-prefix/prefix-length
      — ipv6-prefix-list ipv6-prefix-list-name [create]
      — no ipv6-prefix-list ipv6-prefix-list-name
        — description description-string
        — no description
        — [no] prefix ipv6-prefix/prefix-length
```

3.1.1.2 Show Commands

```
show
  — qos
    — match-list
      — ip-prefix-list
      — ip-prefix-list prefix-list-name [association]
      — ipv6-prefix-list
      — ipv6-prefix-list prefix-list-name [association]
```

3.1.2 Command Descriptions

- [Configuration Commands](#)
- [Show Commands](#)

3.1.2.1 Configuration Commands

match-list

Syntax	match-list
Context	config>qos
Description	This command enables the context to display information for match lists used in QoS policies.

ip-prefix-list

Syntax	ip-prefix-list <i>ip-prefix-list-name</i> [create] no ip-prefix-list <i>ip-prefix-list-name</i>
Context	config>qos>match-list
Description	This command creates a list of IPv4 prefixes for match criteria in QoS policies. An IP prefix list must contain only IPv4 address prefixes created using the prefix command and cannot be deleted if it is referenced by a QoS policy. The no form of this command deletes the specified list.
Parameters	<i>ip-prefix-list-name</i> — A string of up to 32 characters of printable ASCII characters. If the string contains special characters (#, \$, spaces, etc), the entire string must be enclosed within double quotes. The name default (case insensitive) is reserved by the system.

description

Syntax	description <i>description-string</i> no description
Context	config>qos>match-list>ip-prefix-list config>qos>match-list>ipv6-prefix-list

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.
Parameters	<i>description-string</i> — 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.

ipv6-prefix-list

Syntax	ipv6-prefix-list <i>ipv6-prefix-list-name</i> [create] no ipv6-prefix-list <i>ipv6-prefix-list-name</i>
Context	config>qos>match-list
Description	This command creates a list of IPv6 prefixes for match criteria in QoS policies. An ipv6-prefix-list must contain only IPv6 address prefixes created using the prefix command and cannot be deleted if it is referenced by a QoS policy. The no form of this command deletes the specified list.
Parameters	<i>ipv6-prefix-list-name</i> — A string of up to 32 characters of printable ASCII characters. If special characters are used (#, \$, spaces, etc), the string must be enclosed within double quotes. The name default (case insensitive) is reserved by the system. create — Creates IPv6 prefixes for match criteria in QoS policies.

prefix

Syntax	[no] prefix <i>ip-prefix/prefix-length</i>
Context	config>qos>match-list>ip-prefix-list
Description	This command adds an IPv4 address prefix to an existing IPv4 address prefix match list. To add a set of unique prefixes, execute the command with all unique prefixes. The prefixes are allowed to overlap IPv4 address space. An IPv4 prefix addition will be blocked, if resource exhaustion is detected anywhere in the system because of QoS Policies that use this IPv4 address prefix list. The no form of this command deletes the specified prefix from the list.

Parameters	<i>ip-prefix</i> — A valid IPv4 address prefix in dotted decimal notation.
Values	0.0.0.0 to 255.255.255.255 (host bit must be 0)
	<i>prefix-length</i> — Length of the entered IP prefix
Values	1 to 32

prefix

Syntax	[no] prefix <i>ipv6-prefix/prefix-length</i>
Context	config>qos>match-list>ipv6-prefix-list
Description	This command adds an IPv6 address prefix to an existing IPv6 address prefix match list. To add set of unique prefixes, execute the command with all unique prefixes. The prefixes are allowed to overlap IPv6 address space. An IPv6 prefix addition will be blocked if resource exhaustion is detected anywhere in the system because of QoS Policies that use this IPv6 address prefix list. The no form of this command deletes the specified prefix from the list.
Parameters	<i>ipv6-prefix</i> — Specifies the IPv6 prefix for the IP match criterion in hex digits. Values <i>ipv6-address</i>: x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D <i>prefix-length</i> — Specifies the IPv6 prefix length for the IPv6 address expressed as a decimal integer. Values 1 to 128

3.1.2.2 Show Commands

ip-prefix-list

Syntax	ip-prefix-list ip-prefix-list <i>prefix-list-name</i> [association]
Context	show>qos>match-list
Description	Displays the list of configured IPv4 QoS prefix lists or the details of a specific IPv4 QoS prefix list together with the SAP and network QoS policies in which it is used and the entry number within that policy.

Parameters *prefix-list-name* — Specifies an IPv4 prefix list which contains IPv4 address prefixes to be matched.

Values A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

association — Displays the QoS policy and entry number in which the specified prefix list is used.

Output The following output is an example of ip-prefix-list information

Sample Output

```
*A:PE# show qos match-list ip-prefix-list

=====
QoS Match IP Prefix List
=====
Prefix List Name          Description                      Num Prefixes
-----
ipv4list1                  1
ipv4list2                  1
-----
No. of Prefix-List: 2
=====
*A:PE#

*A:PE# show qos match-list ip-prefix-list "ipv4list1"

=====
QoS Match IP Prefix List
=====
Prefix Name      : ipv4list1
Description      : (Not Specified)
-----
IP Prefixes
-----
10.0.0.0/8
-----
No. of Prefixes : 1
-----

-----
Association
-----
QoS Policy ID          Criteria Entry
-----
10          (sap-ingress)    10          (source-ip)
10          (sap-egress)    10          (source-ip)
-----
=====
*A:PE#
```

ipv6-prefix-list

Syntax	ipv6-prefix-list ipv6-prefix-list <i>prefix-list-name</i> [association]
Context	show>qos>match-list
Description	Displays the list of configured IPv6 QoS prefix lists or the details of a specific IPv6 QoS prefix list together with the network QoS policies in which it is used and the entry number within that policy.
Parameters	<i>prefix-list-name</i> — Specifies an IPv6 prefix list which contains IPv6 address prefixes to be matched. Values A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes. association — Displays the QoS policy and entry number in which the specified prefix list is used.
Output	The following output is an example of ipv6-prefix-list information

Sample Output

```
*A:PE# show qos match-list ipv6-prefix-list

=====
QoS Match IPv6 Prefix List
=====
Prefix List Name          Description          Num Prefixes
-----
ipv6list1                  1
ipv6list2                  1
-----
No. of Prefix-List: 2
=====
*A:PE#

*A:PE# show qos match-list ip-prefix-list "ipv6list1"

=====
QoS Match IPv6 Prefix List
=====
Prefix Name      : ipv6list1
Description      : (Not Specified)
-----
IP Prefixes
-----
2001:db8::/32
-----
No. of Prefixes : 1
=====
*A:PE#
```



```
*A:PE# show qos match-list ipv6-prefix-list "ipv6list1" association
```

```
=====
QoS Match IPv6 Prefix List
=====
Prefix Name      : ipv6list1
Description      : (Not Specified)
-----
IPv6 Prefixes
-----
2001:db8::/32
-----
No. of Prefixes : 1
-----

-----
Association
-----
QoS Policy ID      Criteria Entry
-----
10                  (sap-ingress)    10                  (source-ip)
-----
=====
*A:PE#
```

4 Network QoS Policies

Network QoS policies can be applied to network interfaces, CSC network interfaces in a VPRN, pseudowires in VLL and VPLS services, and IES or VPRN spoke interfaces to provide ingress and egress QoS control for the traffic on those objects. Network QoS policies can also be applied to provide aggregate ingress QoS on VPRN and VXLAN services.

4.1 Network QoS Policies Overview

The ingress component of a network QoS policy defines the packet classification 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 defaults to that defined in the default network QoS policy until an explicit policy is defined for the network interface. Packets can also be redirected to an ingress FP queue group.

The egress component of a network QoS policy allows further egress classification to internal forwarding class and profile state. The egress of the network QoS policy also defines the marking based on the forwarding class and the profile state. Packets can also be redirected to an egress port queue group.

Network policy 1 is applied to all network interfaces by default. It cannot be modified or deleted. It defines the default DSCP-to-FC mapping and MPLS EXP-to-FC mapping for the ingress. For the egress, it defines eight forwarding classes and the packet marking criteria.

New (non-default) network policy parameters can be modified. 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 networks to which the policy is applied. For this reason, when a policy requires several changes, it is recommended to copy the policy to a work area policy-id. This work-in-progress copy can be modified until all the changes are made, then the original policy can be overwritten using the **config qos copy** command.

4.1.1 Storing Match-Criteria Entries

FP2-, FP3-, and FP4-based cards store QoS policy match-criteria entries in dedicated memory banks in hardware, also referred to as CAM tables:

- IP/MAC ingress
- IP/MAC egress
- IPv6 ingress
- IPv6 egress

4.1.1.1 FP4-based Cards

To optimize both scale and performance, policy match-criteria entries configured by the operator are compressed by each FP4 line card prior to being installed in hardware.

This compression can result, in an unexpected scenario typically only achieved in a lab environment, in an overload condition where entries for a line card QoS policy on a line card FP are not programmed. This overload condition can occur when applying a QoS policy for the first time on a line card FP or when adding entries to a QoS policy.

Applying a QoS Policy

A policy is installed for the first time on a line card FP if no router interface, service interface, SAP, spoke SDP, mesh SDP or ESM subscriber host was using the policy on this FP.

A policy installed for the first time on a line card FP can lead to a compression failure resulting in an overload condition for this policy on this FP. In this case, none of the entries for the affected QoS match-criteria policy are programmed, and the QoS policy queuing, FC mapping, and FC marking are unaffected.

Adding QoS Match-Criteria Entries

Adding an additional entry to a QoS policy can lead to a compression failure resulting in an overload condition.

In this case, the newly added entry is not programmed on the affected FP. Additional entries added to the same policy after the first overload condition are also not programmed on the affected FP as the system attempts to install all outstanding additions in order.

The CPM QoS management task controls the maximum number of match criteria entries per FP. If the operator attempts to go over the scaling limit, the system will return an interactive error message.

Note: A trap is raised if a policy is in overload, there is no interactive error message.

Removing QoS Match-Criteria Entries

Removing match-criteria entries from a QoS policy is always successful.

Resolving Overload

The overload condition should be resolved by the network operator before adding new entries in the affected policy.

To identify the affected policy, the system logs the overload event providing slot number, FP number, and impacted memory bank. Based on this information, the **tools>dump>qos>match-criteria-overload** command allows the operator to identify the affected policy and policy entries in the system.

To resolve the overload condition, the network operator can remove the newly added entries from the affected policy or assign a different policy.

4.2 Network Ingress

The following types of QoS mapping decisions are applicable to the ingress of a network QoS policy:

- Default QoS mapping
- Ethernet dot1p value mapping (if defined)
- IP DSCP mapping

The inner DSCP marking can be used at the ingress of an LER (see [Network Ingress Tunnel QoS Override](#))

- IP criteria mapping
- IPv6 criteria mapping
- MPLS LSP EXP mapping

The default QoS mapping always exists and every received packet will be mapped to this default if no explicitly defined matching entry exists.

The packets for a specified forwarding class can be redirected to an ingress FP queue group.

4.2.1 Network Ingress Tunnel QoS Override

By default, 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.

Tunnel termination QoS override only applies to IP routing decisions when the tunnel encapsulation is removed. Non-IP routed packets within a terminating tunnel are ignored by the override and are forwarded as described in [Network Ingress](#).

Any tunnel received on the ingress IP interface that traverses the node (where the node is not the ultimate hop for the tunnel) is not affected by the QoS override mechanism and is forwarded as described in [Network Ingress](#).

Tunnel termination QoS override, 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 useful 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 using the ingress **ler-use-dscp** parameter, the IP interface will ignore the tunnel's QoS mapping and derive the internal forwarding class and profile based on the precedence or DSCP values within the routed IP header ToS field compared to the network QoS policy defined on the IP interface.

4.2.2 Network Ingress IP Match Criteria

IP match criteria classification is supported in the ingress section of a network QoS policy.

The classification only applies to the outer IPv4 header of non-tunneled traffic. Consequently, the use of an IP criteria statement in a network QoS policy is ignored for received traffic when the network QoS policy is applied on the ingress network IP interface in the following cases:

- Mesh SDPs in VPLS services
- Spoke SDPs in VPLS and Xpipe services
- Spoke SDPs under an IP interface in an IES or VPRN service
- Spoke SDPs in a VPRN service
- Bindings created automatically by the **auto-bind-tunnel** command in a VPRN service
- IPv6 over IPv4 tunnels
- VXLAN bindings (egress VTEP, VNI)

The only exception is for traffic received on a Draft Rosen tunnel, for which classification on the outer IP header only is supported.

Attempting to apply a network QoS policy containing an IP criteria statement to any object except a network IP interface will result in an error.

The following is an example configuration:

```
configure
  qos
    network 10 name "10" create
    ingress
      ip-criteria
        entry 10 create
        match
```

```

                                dst-ip 10.0.0.1/32
                                exit
                                action fc "h2" profile in
                                exit
                                exit
                                exit
                                exit
                                exit
                                exit

```

4.2.3 Network Ingress IPv6 Match Criteria

IPv6 match criteria classification is supported in the ingress section of a network QoS policy.

The classification only applies to the outer IPv6 header of non-tunneled traffic; consequently, the use of an `ipv6-criteria` statement in a network QoS policy is ignored for received traffic when the network QoS policy is applied on the ingress network IP interface in the following cases:

- Mesh SDPs in VPLS services
- Spoke SDPs in VPLS and Xpipe services
- Spoke SDPs under an IP interface in an IES or VPRN service
- Spoke SDPs in a VPRN service
- Bindings created automatically by the **auto-bind-tunnel** command in a VPRN service
- IPv6 over IPv4 tunnels
- VXLAN bindings (egress VTEP, VNI)

Attempting to apply a network QoS policy containing an IPv6 criteria statement to any object except a network IP interface will result in an error.

The following is an example configuration:

```

configure
  qos
    network 10 name "10" create
      ingress
        ipv6-criteria
          entry 10 create
            match
              dst-ip 2001:db8:1000::1/128
            exit
            action fc "ef" profile in
          exit
        exit
      exit
    exit
  exit

```

exit

4.3 Network Egress

The following types of QoS mapping decisions are applicable to the egress of a network QoS policy:

- IP DSCP mapping
- IP ToS precedence mapping
- IP criteria mapping
- IPv6 criteria mapping

Default dot1p/DE, DSCP, and EXP packet marking are defined for each forwarding class for in and out of profile packets, together with the remarking capability based on the trusted state of the packet's ingress interface.

The packets of a specified forwarding class can be redirected to an egress port queue group.

4.3.1 Egress Packet Reclassification Based on IP precedence DSCP

The user enables IP precedence or DSCP based egress reclassification 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 {exceed | out | in | inplus}]
```

```
config>qos>network>egress>dscp dscp-name [fc fc-name] [profile {exceed | out | in | inplus}]
```

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.

The IP precedence and DSCP based re-classification are supported on a network interface, on a CSC network interface in a VPRN, and 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 an L2 spoke-SDP.

In addition, the egress re-classification commands only take effect if the redirection of the spoke-SDP or CSC interface to use an egress port queue-group succeeds; for example, the following CLI commands succeed:

```
config>service>vprn>if>spoke-sdp>egress>qos network-policy-id port-redirect-group queue-group-name instance instance-id
```

```
config>service>ies>if>spoke-sdp>egress>qos network-policy-id port-redirect-group queue-group-name instance instance-id
```

```
config>service>vprn>nw-if>qos network-policy-id port-redirect-group queue-group-name instance instance-id
```

When the redirection command fails in CLI, the PW uses the network QoS policy assigned to the network IP interface, however any reclassification in the network QoS policy applied to the network interface will be ignored.

4.3.2 Network Egress IP Match Criteria

IP match criteria classification is supported in the egress section of a network QoS policy.

The configuration of egress prec/DSCP classification and the configuration of an egress IP criteria entry statement within a network QoS policy are mutually exclusive.

The criteria action statement port redirect group is not supported on the 7750 SR-a4/a8.

Network QoS policies containing egress IP criteria entry statements are only applicable to network interfaces.

Network egress IP match criteria classification is not supported on an HSMDA.

The following is an example configuration:

```
configure
  qos
    network 10 name "10" create
    egress
      ip-criteria
        entry 10 create
        match
          dst-ip 192.168.1.0/24
        exit
        action fc "af" profile out
      exit
    exit
  exit
exit
```

4.3.3 Network Egress IPv6 Match Criteria

IPv6 match criteria classification is supported in the egress section of a network QoS policy.

The configuration of egress prec/DSCP classification and the configuration of an egress IPv6 criteria entry statement within a network QoS policy are mutually exclusive.

The criteria action statement port redirect group is not supported on the 7750 SR-a4/a8.

Network QoS policies containing egress IPv6 criteria entry statements are only applicable to network interfaces.

Network egress IPv6 match criteria classification is not supported on an HSMDA.

The following is an example configuration:

```
configure
  qos
    network 10 name "10" create
    egress
      ipv6-criteria
        entry 10 create
        match
          dst-ip 2001:db8:2000::1/128
        exit
      action fc "ef" profile in
    exit
  exit
exit
```

4.4 QoS for Self-Generated (CPU) Traffic on Network Interfaces

Differentiated services code point (DSCP), forwarding class (FC), and IEEE 802.1p values can be specified to be used by protocol packets generated by the node. This enables prioritization or deprioritization of every protocol (as required). The markings effect a change in behavior on ingress when queuing.

DSCP marking for internally generated control and management traffic 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 than for a VPRN service. ARP, IS-IS, and PPPoE are not IP protocols, so only 802.1p values can be configured.

The DSCP value can be set per application. When an application is configured to use a specified DSCP value, the 802.1p and MPLS EXP bits will be marked in accordance with the network (default 802.1p value of 7) or access (default 802.1p value of 0) egress policy as it applies to the logical interface that the packet will be egressing.

The configuration of self-generated QoS is supported in the base router, VPRN, and management contexts.

The default values for self-generated traffic on network interfaces 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: according to the egress QoS policy (7 by default)
- Management protocols (SSH, SNMP, etc.)
 - Forwarding class: Network Control (NC)
 - DSCP value: AF41
 - 802.1p value: according to the egress QoS policy (7 by default)

The default QoS values for self-generated traffic on network interfaces are listed in [Table 16](#).

Table 16 Default QoS Values for Self-Generated Traffic

Protocol	DSCP
ANCP	NC1

Table 16 Default QoS Values for Self-Generated Traffic (Continued)

Protocol	DSCP
APS	NC1
ARP	N/A
BFD	NC1
BGP	NC1
BMP	AF41
Cflowd	NC1
DHCP	NC1, AF41, NC2
Diameter	AF41
DNS	AF41
FTP	AF41
gRPC	AF41
GTP	NC1, NC2
HTTP	AF41
ICMP	BE, NC1
IGMP	NC1
IGMP Reporter	NC1
IS-IS	N/A
L2TP	NC1
LDP	NC1
LMP	NC1
MLD	NC1
MSC (Multichassis Support)	NC1
MSDP	NC1
Mtrace2	NC1
ND (NDIS)	NC1, NC2
NTP/SNTP	NC1
OpenFlow	NC1

Table 16 Default QoS Values for Self-Generated Traffic (Continued)

Protocol	DSCP
OSPF	NC1
PCEP	NC1
PIM	NC1
PPPoE	N/A
PTP	NC1
RADIUS	NC1
RIP	NC1
RSVP	NC1
sFlow	NC1
SNMP Gets/Sets	AF41
SNMP Traps	AF41
SRRP	NC1
SSH, SCP, SFTP	AF41
Syslog	AF41
TACACS+	AF41
Telnet	AF41
TFTP	AF41
Traceroute	BE
TWAMP, TWAMP Light	N/A
VRRP	NC1
WSC	NC1
XMPP	NC1



Note: ICMP echo requests (type 8) and ICMPv6 echo requests (type 128) initiated from the router will use the DSCP value set by the **sgt-qos** command. The FC value is NC by default, or the value specified in the **ping** command parameter **fc fc-name**.



Note: The DSCP values for TWAMP and TWAMP Light test packets are not configured with **sgt-qos** commands. The DSCP value for TWAMP test packets reflected by the TWAMP server is specified in the TWAMP control process. The DSCP value for TWAMP Light test packets is set by the test configuration. The TWAMP Light reflector uses the arriving TWAMP test packet to determine the return DSCP value.



Note: Some applications have multiple DSCP default values depending on the context or service.



Note: Configurable values for ANCP, APS, BFD, LMP, MSC, OpenFlow, WSC, and XMPP are not supported.

4.4.1 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.

4.5 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.

4.5.1 Creating 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 a 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.
- Egress criteria can be modified 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. The forwarding class criteria define the egress characteristics of the queue and the marking criteria of packets flowing through it.
 - DE marking — Specifies that the DE bit should be marked based on whether the packet profile is in-profile, inplus-profile, out-of-profile, or exceed-profile.
 - dot1p — The dot1p value is used for all VLAN-tagged packets requiring marking that egress on this forwarding class queue, with the option of specifying a different value for packets that are in-profile or out-of-profile. Inplus-profile traffic is marked with the same values as in-profile traffic. Exceed-profile traffic is marked with the same values as out-of-profile traffic.

- DSCP — The DSCP value is used for all IP packets requiring marking that egress on this forwarding class queue that are in-profile or out-of-profile. Inplus-profile traffic is marked with the same values as in-profile traffic. Exceed-profile traffic is marked with the same values as out-of-profile traffic.
- LSP EXP — The EXP value is used for all MPLS-labeled packets requiring marking that egress on this forwarding class queue that are in-profile or out-of-profile. Inplus-profile traffic is marked with the same values as in-profile traffic. Exceed-profile traffic is marked with the same values as out-of-profile traffic.
- Port redirection — Specifies that the traffic should be redirected to a network egress queue group policer or queue.
- DSCP — Creates a mapping between the DSCP of the network egress traffic and the forwarding class and profile. Egress traffic that matches the specified DSCP is assigned to the corresponding forwarding class with the specified profile.
- IP criteria — Creates a mapping between the possible match criteria of the network egress traffic and the forwarding class and profile. Egress traffic that matches the IPv4 criteria is assigned to the corresponding forwarding class and profile.
- IPv6 criteria — Creates a mapping between the possible match criteria of the network egress traffic and the forwarding class and profile. Egress traffic that matches the IPv6 criteria is assigned to the corresponding forwarding class and profile.
- Prec — Creates a mapping between the IP precedence of the network egress traffic and the forwarding class and profile. Egress traffic that matches the specified IP precedence is assigned to the corresponding forwarding class with the specified profile.
- Ingress criteria — Specifies the DSCP- or dot1p-to-forwarding class mapping for all IP packets and defines 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 configured classification. The default action specifies the forwarding class and profile to which such packets are assigned.
 - dot1p — Creates a mapping between the dot1p of the network ingress traffic and the forwarding class and profile. Ingress traffic that matches the specified dot1p is assigned to the corresponding forwarding class and profile.
 - DSCP — Creates a mapping between the DSCP of the network ingress traffic and the forwarding class and profile. Ingress traffic that matches the specified DSCP is assigned to the corresponding forwarding class.

- Forwarding class — The forwarding class name represents an ingress queue.
- IP criteria — Creates a mapping between the possible match criteria of the network ingress traffic and the forwarding class and profile. Ingress traffic that matches the IPv4 criteria is assigned to the corresponding forwarding class and profile.
- IPv6 criteria — Creates a mapping between the possible match criteria of the network ingress traffic and the forwarding class and profile. Ingress traffic that matches the IPv6 criteria is assigned to the corresponding forwarding class and profile.
- LER use DSCP — Specifies that DSCP matching based on the tunneled IP packet should be used on an LER rather than matching on the outer encapsulated header.
- LSP EXP — Creates a mapping between the LSP EXP bits of the network ingress traffic and the forwarding class and profile. Ingress traffic that matches the specified LSP EXP bits is assigned to the corresponding forwarding class and profile.

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

CLI Syntax:

```

config>qos# network network-policy-id [create]
      description description-string
      egress
        dscp dscp-name fc fc-name profile {in | out |
          exceed | inplus}
        fc {be | l2 | af | l1 | h2 | ef | h1 | nc}
          de-mark [force de-value]
          dot1p dot1p-priority
          dot1p-in-profile dot1p-priority
          dot1p-out-profile dot1p-priority
          dscp-in-profile dscp-name
          dscp-out-profile dscp-name
          lsp-exp-in-profile lsp-exp-value
          lsp-exp-out-profile lsp-exp-value
          port-redirect-group {queue queue-id |
            policer plcr-id [queue queue-id]}
        prec dscp-name fc fc-name profile {in | out |
          exceed | inplus}
          remarking
        ip-criteria
          entry entry-id [create]
            action [fc fc-name profile {in | out
              | exceed | inplus}]
              [port-redirect-group {queue
                queue-id | policer policer-id
                [queue queue-id]}]

```

```

description description-string
match [protocol protocol-id]
    dscp dscp-name
    dst-ip {ip-address/mask | ip-
        address ipv4-address-mask}
    dst-port [{lt | gt | eq} {dst-
        port-number | range start
        end}]
    fragment {true | false}
    icmp-type icmp-type
    src-ip {ip-address/mask | ip-
        address ipv4-address-mask}
    src-port [{lt | gt | eq} {src-
        port-number | range start
        end}]
    renum old-entry-id new-entry-id
ipv6-criteria
entry entry-id [create]
    action fc fc-name profile {in | out
    | exceed | inplus}]
    [port-redirect-group {queue
        queue-id | policer policer-id
        [queue queue-id]}]
description description-string
match [next-header next-header]
    dscp dscp-name
    dst-ip {ipv6-address/mask |
        ipv6-address ipv6-address-
        mask}
    dst-port [{lt | gt | eq} {dst-
        port-number | range start
        end}]
    fragment {true | false | first-
        only | non-first-only}
    src-ip {ipv6-address/mask |
        ipv6-address ipv6-address-
        mask}
    src-port [{lt | gt | eq} src-
        port-number
        src-port range start end
    renum old-entry-id new-entry-id
ler-use-dscp
lsp-exp lsp-exp-value fc fc-name profile {in |
    out}
scope {exclusive | template}

```

```

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#

```

4.5.2 Applying Network QoS Policies

Use the following CLI syntax to apply network policies to the router access uplink port's IP 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#

```

4.5.3 Default Network QoS Policy Values

The default network policy for IP interfaces is identified as policy-id 1. Default policies cannot be modified or deleted. [Table 17](#) lists default network policy parameters.

Table 17 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
			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

Table 17 Network Policy Defaults (Continued)

Field				Default	
			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	
		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				

Table 17 Network Policy Defaults (Continued)

Field			Default	
		dscp-in-profile	af41	
		dscp-out-profile	af42	
		lsp-exp-in-profile	4	
		lsp-exp-out-profile	4	
	fc l			
		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	
		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 11
        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 12
        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#
```

4.6 Service Management Tasks

4.6.1 Deleting QoS Policies

A network policy is associated by default with router interfaces.

The default policy can be replaced with a non-default policy, but cannot be removed from the configuration. When a non-default policy is removed, 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
            unreachable 100 10
            ttl-expired 100 10
        exit
        qos 1
        ingress
            no filter
        exit
        egress
            no filter
        exit
        no mac
        no cflowd
        no shutdown
    exit
    interface "ALA-1-3"
    ...
#-----
A:ALA-7>config>router#
```

4.6.2 Removing a Policy from the QoS Configuration

To delete a network policy, enter the following commands:

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

4.6.3 Copying and Overwriting Network Policies

An existing network policy can be copied 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#
```

4.6.4 Editing QoS Policies

Existing policies, except the default policies and entries in the CLI, can be changed. 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 in a work area, make the edits, then overwrite the original policy.

4.7 Network QoS Policy Command Reference

4.7.1 Command Hierarchies

- Configuration Commands
 - Multi-Class Frame Relay Profile Commands
 - Network QoS Policy Commands
 - Model-Driven Automatic ID Commands
 - Self-Generated Traffic Commands
- Operational Commands
- Show Commands
- Debug Commands

4.7.1.1 Configuration Commands

4.7.1.1.1 Multi-Class Frame Relay Profile Commands

```
config
  — qos
    — mc-fr-profile-ingress profile-id [create]
    — no mc-fr-profile-ingress profile-id
      — class class-id
        — reassembly-timeout timeout-value
        — no reassembly-timeout
      — description description-string
      — no description
    — mc-fr-profile-egress profile-id [create]
    — no mc-fr-profile-egress profile-id
      — class class-id
        — max-queue-size queue-size
        — no max-queue-size
        — mir mir
        — no mir
        — weight weight
        — no weight
      — description description-string
      — no description
```

4.7.1.1.2 Network QoS Policy Commands

The commands listed in this section apply only to the 7450 ESS and 7750 SR.

```

config
  — qos
    — network network-policy-id [create] [name name]
    — no network network-policy-id
      — description description-string
      — no description
      — egress
        — dscp dscp-name fc fc-name profile {in | out | exceed | inplus}
        — no dscp dscp-name
        — fc fc-name
        — no fc
          — 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
          — port-redirect-group {queue queue-id | policer policer-id
            [queue queue-id]}
          — no port-redirect-group
        — [no] ip-criteria
          — entry entry-id [create]
          — no entry entry-id
            — action [fc fc-name profile {in | out | exceed | inplus}
              [port-redirect-group {queue queue-id | policer
                policer-id [queue queue-id]}]
            — 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 | eq} dst-port-number
              — dst-port range start end
              — no dst-port

```


- **fragment** {true | false}
- **no fragment**
- **icmp-type** *icmp-type*
- **no icmp-type**
- **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-port**
- **renum** *old-entry-id new-entry-id*
- [no] **ipv6-criteria**
 - **entry** *entry-id* [create]
 - **no entry** *entry-id*
 - **action** [fc *fc-name* **profile** {in | out | exceed | inplus}] [**port-redirect-group** {*queue queue-id* | **policer** *policer-id* [*queue queue-id*]}]
 - **no action**
 - **description** *description-string*
 - **no description**
 - **match** [*next-header next-header*]
 - **no match**
 - **dscp** *dscp-name*
 - **no dscp**
 - **dst-ip** {*ipv6-address/mask* | *ipv6-address* *ipv6-address-mask* | **ipv6-prefix-list** *ipv6-prefix-list-name*}
 - **no dst-ip**
 - **dst-port** {lt | gt | eq} *dst-port-number*
 - **dst-port** range *start end*
 - **no dst-port**
 - **fragment** {true | false | first-only | non-first-only}
 - **no fragment**
 - **icmp-type** *icmp-type*
 - **no icmp-type**
 - **src-ip** {*ipv6-address/mask* | *ipv6-address* *ipv6-address-mask* | **ipv6-prefix-list** *ipv6-prefix-list-name*}}
 - **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*
 - **prec** *ip-prec-value* **fc** *fc-name* **profile** {in | out | exceed | inplus}
 - **no prec** *ip-prec-value*
 - **remarking** [force]
 - **no remarking**
 - **ingress**
 - **default-action** **fc** *fc-name* **profile** {in | out}
 - **dot1p** *dot1p-priority* **fc** *fc-name* **profile** {in | out | use-de}
 - **no dot1p**
 - **dscp** *dscp-name* **fc** *fc-name* **profile** {in | out}
 - **no dscp**

```

— fc fc-name
— no fc
  — fp-redirect-group policer-type policer-id
  — no fp-redirect-group policer-type
— [no] ip-criteria
  — entry entry-id [create]
  — 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 ipv4-address-mask | ip-prefix-list ip-prefix-list-name}
      — 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/mask | ip-address ipv4-address-mask | ip-prefix-list ip-prefix-list-name}
      — 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
  — entry entry-id [create]
  — no entry entry-id
    — action 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 | ipv6-address ipv6-address-mask | ipv6-prefix-list ipv6-prefix-list-name}
      — no dst-ip
      — dst-port {lt | gt | eq} dst-port-number
      — dst-port range start end
      — no dst-port
      — fragment {true | false | first-only | non-first-only}
      — no fragment

```

- **src-ip** {*ipv6-address/prefix-length* | *ipv6-address*
ipv6-address-mask | **ipv6-prefix-list** *ipv6-*
prefix-list-name}
- **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] ler-use-dscp**
- **lsp-exp** *lsp-exp-value fc fc-name profile {in | out}*
- **no lsp-exp**
- **scope** {*exclusive* | *template*}
- **no scope**

4.7.1.1.3 Model-Driven Automatic ID Commands

- ```
config
 — qos
 — md-auto-id
 — qos-policy-id-range start policy-id end policy-id
 — no qos-policy-id-range
```

#### 4.7.1.2 Self-Generated Traffic Commands

- ```
config
  — service
    — vprn
      — 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
      — 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
```

4.7.1.3 Operational Commands

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

#### 4.7.1.4 Show Commands

```

show
 — qos
 — dscp-table [value dscp-value]
 — mc-fr-profile-ingress
 — mc-fr-profile-ingress profile-id [detail]
 — mc-fr-profile-egress
 — mc-fr-profile-egress profile-id [detail]
 — md-auto-id
 — network [network-policy-id] [detail] [match-criteria]
 — router
 — sgt-qos
 — application [app-name] [dscp | dot1p]
 — dscp-map [dscp-name]

```

#### 4.7.1.5 Debug Commands

```

tools
 — dump
 — qos
 — match-criteria-overload
 — match-criteria-overload card slot-number fp fp-number
 — match-criteria-overload card slot-number fp fp-number network network-
 policy-id direction direction
 — match-criteria-overload card slot-number fp fp-number sap-egress policy-id
 — match-criteria-overload card slot-number fp fp-number sap-ingress policy-id

```

### 4.7.2 Command Descriptions

- Configuration Commands
  - Generic Commands
  - Multi-Class Frame Relay Profile Commands
  - Network QoS Policy Commands
    - Network Ingress QoS Policy Commands
    - Network Egress QoS Policy Commands
    - Network Ingress QoS Policy Forwarding Class Commands
    - Network Egress QoS Policy Forwarding Class Commands
    - Network Ingress and Egress Criteria QoS Policy Commands

- [Model-Driven Automatic ID Commands](#)
- [Self-Generated Traffic Commands](#)
- [Operational Commands](#)
- [Show Commands](#)

## 4.7.2.1 Configuration Commands

### 4.7.2.1.1 Generic Commands

#### description

|                    |                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>description</b> <i>description-string</i><br><b>no description</b>                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>qos>mc-fr-profile-egress<br>config>qos>mc-fr-profile-ingress<br>config>qos>network<br>config>qos>network>egress>ip-criteria>entry<br>config>qos>network>egress>ipv6-criteria>entry<br>config>qos>network>ingress>ip-criteria>entry<br>config>qos>network>ingress>ipv6-criteria>entry                                                          |
| <b>Description</b> | This command creates a text description stored in the configuration file for a configuration context.<br><br>The <b>description</b> command associates a text string with a configuration context to help identify the context in the configuration file.<br><br>The <b>no</b> form of this command removes any description string from the context. |
| <b>Parameters</b>  | <i>description-string</i> — A text string describing the entity. Allowed values are any string of up to 80 printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.                                                                              |

### 4.7.2.1.2 Multi-Class Frame Relay Profile Commands

#### mc-fr-profile-ingress

**Syntax**    **mc-fr-profile-ingress** *profile-id* [**create**]

---

**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.

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

**Values** 1 to 65535

**create** — Required parameter when creating a multi-class Frame Relay ingress profile.

## class

**Syntax** **class** *class-id*

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

**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.

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

**Values** 0 to 3

## reassemble-timeout

**Syntax** **reassemble-timeout** *timeout-value*  
**no reassemble-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** reassemble-timeout 25 — class 0  
reassemble-timeout 25 — class 1  
reassemble-timeout 100 — class 2  
reassemble-timeout 1000 — class 3

---

**Parameters**    *timeout-value* — Specifies the timeout value, in ms.  
**Values**        1 to 1000

## mc-fr-profile-egress

**Syntax**        **mc-fr-profile-egress** *profile-id* [**create**]  
                   **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 profiles may be created on the system.

                  The no form of this command deletes the profile.

**Parameters**    *profile-id* — Specifies the profile number.  
**Values**        1 to 65535  
**create** — Required parameter when creating a multi-class Frame Relay egress profile.

## 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**        max-queue-size 10 — class 0  
                   max-queue-size 50 — class 1  
                   max-queue-size 150 — class 2  
                   max-queue-size 750 — class 3

**Parameters**    *queue-size* — Specifies the number, in ms, of the available link or bundle rate.  
**Values**        1 to 1000

## mir

**Syntax**        **mir** *mir*

**no mir**

|                    |                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>qos>mc-fr-profile-egress>class                                                                                                   |
| <b>Description</b> | This command configures the minimum information rate scheduling parameter for each Frame Relay scheduling class queue for this profile. |
| <b>Default</b>     | 90% for all classes                                                                                                                     |
| <b>Parameters</b>  | <i>mir</i> — Specifies the percentage of the available link or bundle rate.                                                             |
| <b>Values</b>      | 1 to 100                                                                                                                                |

## weight

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>weight</b> <i>weight</i><br><b>no weight</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>qos>mc-fr-profile-egress>class                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command configures the WRR weight scheduling parameter for each Frame Relay scheduling class queue for this profile.</p> <p>The weight parameter is not applicable to class 0 because the class 0 queue is serviced by the scheduler with the highest priority all the time over all other class queues.</p> <p>After Class 0 queue packets are serviced, Class 1 queue packets are the next priority up to a rate of MIR. Beyond MIR, Class1, Class2, and Class 3 queues are serviced in proportion to the configured weights.</p> |
| <b>Parameters</b>  | <i>weight</i> — Specifies the weight schedule.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Values</b>      | 1 to 100                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>     | weight 1 — class 1 (not configurable)<br>weight 90 — class 2<br>weight 10 — class 3                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**4.7.2.1.3 Network QoS Policy Commands**

## network

|                |                                                                                                                                     |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>network</b> <i>network-policy-id</i> [ <b>create</b> ] [ <b>name</b> <i>name</i> ]<br><b>no network</b> <i>network-policy-id</i> |
| <b>Context</b> | config>qos                                                                                                                          |



|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command creates or edits a QoS network policy. The network policy defines the treatment that IP or MPLS packets receive as they ingress and egress the network port.</p> <p>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.</p> <p>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 interface access 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, dot1p/DE, or MPLS EXP bit marking based on the forwarding class and the profile state.</p> <p>Network <b>policy-id 1</b> exists as the default policy that is applied to all network interfaces by default. The network <b>policy-id 1</b> cannot be modified or deleted. It defines the default DSCP-to-FC mapping and MPLS EXP-to-FC mapping for the ingress. For the egress, it defines six forwarding classes that represent individual queues and the packet marking criteria.</p> <p>Network <b>policy-id 1</b> exists as the default policy that is applied to all network ports by default. This default policy cannot be modified or deleted. It defines the default DSCP-to-FC mapping and default unicast meters for ingress IP traffic. For the egress, it defines the forwarding class to dot1p and DSCP values and the packet marking criteria.</p> <p>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 <b>ip-interface</b> or the DSCP-to-FC mapping (for network QoS policy of type <b>port</b>). 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). Parameters can be modified or the <b>no</b> form of this command can be used to remove an object from the configuration.</p> <p>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, then written over the original policy-id. Use the config qos copy command to maintain policies in this manner.</p> <p>The <b>no</b> 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 <b>policy policy-id 1</b> cannot be deleted.</p> |
| <b>Default</b>     | network 1 — System Default Network Policy 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

---

|                   |                                                                                                                                                                                                      |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>network-policy-id</i> — The policy-id uniquely identifies the policy on the router.                                                                                                               |
| <b>Values</b>     | 1 to 65535                                                                                                                                                                                           |
| <b>Default</b>    | 1                                                                                                                                                                                                    |
| <b>create</b>     | — Required parameter when creating a QoS network policy.                                                                                                                                             |
| <b>name</b>       | <i>name</i> — A name that is saved as part of the configuration data. If a name is not specified at creation time, then SR OS assigns a string version of the network policy identifier as the name. |
| <b>Values</b>     | A string up to 64 characters                                                                                                                                                                         |

## scope

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>scope {exclusive   template}</b><br><b>no scope</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>qos>network                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | 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.<br><br>The <b>no</b> form of this command sets the scope of the policy to the default of <b>template</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>     | scope template                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <b>exclusive</b> — 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.<br><br>The system default policies cannot be put into the exclusive scope. An error will be generated if the <b>scope exclusive</b> command is executed in any policies with a policy-id equal to 1.<br><br><b>template</b> — When the scope of a policy is defined as template, the policy can be applied to multiple interfaces on the router.<br><br>Default QoS policies are configured with template scopes. An error is generated if the template scope parameter is modified to exclusive scope on default policies. |

## Network Ingress QoS Policy Commands

## ingress

|                |                    |
|----------------|--------------------|
| <b>Syntax</b>  | <b>ingress</b>     |
| <b>Context</b> | config>qos>network |

**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 premarked 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 bit 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-name* — Specifies the forwarding class name. All packets with DSCP value or MPLS EXP or dot1p bits that are not defined will be placed in this forwarding class.

**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-profile 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.

**Values** in, out

## dot1p

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

**Context** config>qos>network>ingress

**Description** This command explicitly sets the forwarding class or enqueueing priority and 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 and be assigned to the forwarding class and enqueueing 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 enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing 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 enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.

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

**Values** 0 to 7

*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-profile or out-of-profile based on this command, or will use the DE bit to determine the profile of the packets (in-profile is used when DE = 0 and out-of-profile is used when DE = 1). In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets.

## dscp

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

**Context** config>qos>network>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 64 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.

**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 lowercase, exactly as shown in the first column in [Table 18](#) and [Table 19](#).

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 18** Default DSCP Names to DSCP Value Mapping

| DSCP Name | DSCP Value<br>Decimal | DSCP Value<br>Hexadecimal | DSCP Value<br>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 19** Default Class Selector Code Points to DSCP Value Mapping

| DSCP Name | DSCP Value<br>Decimal | DSCP Value<br>Hexadecimal | DSCP Value<br>Binary |
|-----------|-----------------------|---------------------------|----------------------|
| cs7       | 56                    | 0x38                      | 0b111000             |
| cs6       | 48                    | 0X30                      | 0b110000             |
| cs5       | 40                    | 0x28                      | 0b101000             |
| cs4       | 32                    | 0x20                      | 0b100000             |
| cs3       | 24                    | 0x18                      | 0b011000             |
| cs2       | 16                    | 0x10                      | 0b010000             |
| cs1       | 08                    | 0x8                       | 0b001000             |

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

**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. 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.

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

DSCP values mapping to forwarding class “be” can only be set to out-of-profile.

**Values** in, out

## 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 that the network-qos-policy-id is associated with. The command may be defined at any time 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 this 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**

**Context** config>qos>network>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** no lsp-exp

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

**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.

**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.

**Values** in, out

---

## Network Egress QoS Policy Commands

### egress

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>egress</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>qos>network                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>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.</p> <p>The forwarding class and profile state mapping to in- and out-of-profile DiffServ Code Points (DSCPs), dot1p, and MPLS EXP bits mapping for all labeled packets are also defined in this context.</p> <p>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.</p> <p>All out-of-profile service packets are marked with the corresponding out-of-profile DSCP, dot1p, or the EXP bit value at network egress. All the in-profile service ingress packets are marked with the corresponding in-profile DSCP, dot1p, or EXP bit value based on the forwarding class to which they belong. The exceed-profile traffic is marked with the same value as out-of-profile traffic and the inplus-profile traffic is marked with the same value as in-profile traffic.</p> |

### dscp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dscp dscp-name fc fc-name profile {in   out   exceed   inplus}</b><br><b>no dscp dscp-name</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>qos>network>egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command configures an IP Differentiated Services Code Point (DSCP) value that must be matched to perform the associated reclassification actions. If an egress packet on an IES/ VPRN interface spoke SDP, on a CSC network interface in a VPRN, or on a network interface that the network QoS policy is applied to, matches the specified IP DSCP value, the forwarding class and profile may be overridden.</p> <p>By default, the forwarding class and profile of the packet are derived from ingress classification and profiling functions. Matching a DHCP-based reclassification rule will override all IP precedence-based reclassification rule actions.</p> <p>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.</p> |



The configuration of egress DSCP classification and the configuration of an egress IP criteria or IPv6 criteria entry statement within a network QoS policy are mutually exclusive.

The IP precedence- and DSCP-based reclassification are supported on a network interface, on a CSC network interface in a VPRN, and on a PW used in an IES or VPRN spoke interface. The CLI will block the application of a network QoS policy with the egress reclassification commands to the spoke SDP part of a Layer 2 service.

Conversely, the CLI will not allow the user to add the egress reclassification commands to a network QoS policy if the policy is being used by a Layer 2 spoke SDP.

The egress reclassification commands will only take effect if the redirection of the spoke SDP or CSC interface to use an egress port queue group succeeds. For example, the following CLI command would be successful:

```
config>service>vprn>if>spoke-sdp>egress>qos network-policy-id port-redirect-group
queue-group-name instance instance-id
```

```
config>service>ies>if>spoke-sdp>egress>qos network-policy-id port-redirect-group
queue-group-name instance instance-id
```

```
config>service>vprn>nw-if>qos network-policy-id port-redirect-group queue-group-
name instance instance-id
```

If the redirection command fails, the PW will use the network QoS policy assigned to the network IP interface, however any reclassification in the network QoS policy applied to the network interface will be ignored.

The **no** form of this command removes the egress reclassification rule.

**Parameters** *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

**fc** *fc-name* — be, l2, af, l1, h2, ef, h1, nc

**profile {in | out | exceed | inplus}** — The profile reclassification action is mandatory.

When specified, packets matching the DSCP value will be explicitly reclassified to the profile specified regardless of the ingress profiling decision. To remove the profile reclassification action for the specified DSCP value, the **no dscp** command must be executed.

**in** - Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

**out** - Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

**exceed** - Specifies that any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

**inplus** - Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

---

## prec

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>prec</b> <i>ip-prec-value</i> <b>fc</b> <i>fc-name</i> <b>profile</b> { <b>in</b>   <b>out</b>   <b>exceed</b>   <b>inplus</b> }<br><b>no prec</b> <i>ip-prec-value</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>network>egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command defines a specific IP precedence value that must be matched in order to perform the associated reclassification actions. If an egress packet on an IES/VPNR interface spoke SDP, on a CSC network interface in a VPRN, or network interface that the network QoS policy is applied to, matches the specified IP precedence value, the forwarding class and profile may be overridden.</p> <p>By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions.</p> <p>The IP precedence bits used to match against the 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, IP precedence-based matching is not performed.</p> <p>The configuration of egress prec classification and the configuration of an egress IP criteria or IPv6 criteria entry statement within a network QoS policy are mutually exclusive.</p> <p>The IP precedence-based and DSCP-based reclassification are supported on a network interface, on a CSC network interface in a VPRN, and on a PW used in an IES or VPRN spoke interface.</p> <p>This command will block the application of a network QoS policy with the egress reclassification commands to a spoke SDP part of a Layer 2 service. Conversely, this command will not allow the user to add the egress reclassification commands to a network QoS policy if it is being used by a Layer 2 spoke SDP.</p> <p>The egress reclassification commands will only take effect if the redirection of the spoke SDP or CSC interface to use an egress port queue-group succeeds. For example, the following commands will succeed:</p> <pre> — config&gt;service&gt;vprn&gt;if&gt;spoke-sdp&gt;egress&gt;qos <i>network-policy-id</i> port-redirect-group   <i>queue-group-name</i> instance <i>instance-id</i> — config&gt;service&gt;ies&gt;if&gt;spoke-sdp&gt;egress&gt;qos <i>network-policy-id</i> port-redirect-group   <i>queue-group-name</i> instance <i>instance-id</i> — config&gt;service&gt;vprn&gt;nw-if&gt; qos <i>network-policy-id</i> port-redirect-group <i>queue-group-name</i>   instance <i>instance-id</i> </pre> <p>When the redirection command fails in CLI, the PW will use the network QoS policy assigned to the network IP interface; however, any reclassification in the network QoS policy applied to the network interface will be ignored.</p> <p>The <b>no</b> form of this command removes the egress reclassification rule.</p> |
| <b>Parameters</b>  | <p><i>ip-prec-value</i> — 0 to 7</p> <p><b>fc</b> <i>fc-name</i> — be, l2, af, l1, h2, ef, h1, nc</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

*profile {in | out | exceed | inplus}* — The profile reclassification action is mandatory. When specified, packets matching the IP precedence value will be explicitly reclassified to the profile specified regardless of the ingress profiling decision. To remove the profile reclassification action for the specified IP precedence value, the **no prec** command must be executed.

This value may be overwritten by an explicit profile action in an DSCP reclassification match.

**in** - Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

**out** - Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

**exceed** - Specifies that any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

**inplus** - Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

## remarking

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>remarking [force]</b><br><b>no remarking</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>qos>network>egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>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.</p> <p>Normally, packets that ingress on network ports have either the DSCP or, for 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 <b>remarking</b> is enabled.</p> <p>Remarking can be required if this router is connected to a different DiffServ domain where the DSCP-to-forwarding class mapping is different.</p> <p>Normally, no remarking is necessary when all router devices are in the same DiffServ domain.</p> <p>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 egress network remark trusted state on each type of ingress IP interface and trust state is listed in <a href="#">Table 20</a>.</p> <p>The remark trusted state has no effect on packets received on an ingress VPRN IP interface.</p> |

**Table 20** Ingress IP Interface Type and Trust State Effect on Egress Network Remarking

| 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 resets the configuration 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.

### Network Ingress QoS Policy Forwarding Class Commands

#### fc

- Syntax** **fc** *fc-name*  
**no fc**
- Context** config>qos>network>ingress
- Description** This command is used to enter the CLI node to configure QoS parameters for the specified forwarding class. The **fc** command overrides the default parameters for that forwarding class from the values defined in the network default policy.
- The **no** form of this command removes the forwarding class name configuration. The forwarding class reverts to the parameters defined in the default network policy.
- Parameters** *fc-name* — The case-sensitive, system-defined forwarding class name for which policy entries will be created.
- Values** be, l2, af, l1, h2, ef, h1, nc

## fp-redirect-group

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fp-redirect-group</b> <i>policer-type</i> <i>policer-id</i><br><b>no fp-redirect-group</b> <i>policer-type</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>qos>network>ingress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command is used to redirect the FC of a broadcast packet received in a VPLS service over a PW or network IP interface to an ingress forwarding plane queue-group.</p> <p>It defines the mapping of an 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 or mesh SDP or a network IP interface.</p> <p>The broadcast-policer statement is ignored when the network QoS policy is applied to any object other than a VPLS spoke or mesh SDP or a network IP interface.</p> <p>The <b>no</b> form of this command removes the redirection of the FC.</p> |
| <b>Parameters</b>  | <p><i>policer-type</i> — The policer type to be used. The <i>policer-type</i> is ignored when the network QoS policy is applied to any object other than a VPLS spoke or mesh SDP or a network IP interface.</p> <p><b>Values</b> broadcast-policer   mcast-policer   policer   unknown-policer</p> <p><i>policer-id</i> — The specified <i>policer-id</i> must exist within the queue-group template applied to the ingress context of the forwarding plane.</p> <p><b>Values</b> 1 to 32</p>                                                                                                                                                                                                                                                                                                       |

## Network Egress QoS Policy Forwarding Class Commands

### fc

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fc</b> <i>fc-name</i><br><b>no fc</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>qos>network>egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command is used to enter the CLI node to configure QoS parameters for the specified forwarding class. The FC name represents a CLI parent node that contains parameters describing the egress marking criteria of packets flowing through it. This command overrides the default parameters for that forwarding class from the values defined in the network default policy. It can also be used to redirect packets to a policer or queue in a network egress queue group instance.</p> <p>The <b>no</b> form of this command removes the forwarding class name configuration. The forwarding class reverts to the parameters defined in the default network policy.</p> |

---

|                   |                                                                                                                     |
|-------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>fc-name</i> — The case-sensitive, system-defined forwarding class name for which policy entries will be created. |
| <b>Values</b>     | be, l2, af, l1, h2, ef, h1, nc                                                                                      |

## de-mark

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>de-mark</b> [ <b>force</b> <i>de-value</i> ]<br><b>no de-mark</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>qos>network>egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command is used to explicitly define the marking of the DE bit for <b>fc</b> <i>fc-name</i> according to the inplus-profile or in-profile and out-of-profile or exceed-profile status of the packet ( <i>fc-name</i> may be used to identify the dot1p value).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><i>de-value</i> — Specifies that this value is to be used for all the packets of this forwarding class regardless of their profile status.</p> <p>If no DE value is present, the default values are used for the marking of the DE bit; that is, 0 for inplus-profile and in-profile packets, 1 for out-of-profile and exceed-profile packets. For more information, refer to the <i>IEEE 802.1ad-2005 standard</i>.</p> <p>In the PBB case, use the following rules for a network port (B-SDP):</p> <ul style="list-style-type: none"> <li>• the outer VID follows the rules for regular SDP</li> <li>• for packets originating 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 is set to zero.</li> <li>• for transit packets (B-SAP/B-SDP to B-SDP), the related ITAG bits are preserved, the same as for BVID.</li> </ul> <p><b>Values</b>      0, 1</p> |

## dot1p

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dot1p</b> <i>dot1p-priority</i><br><b>no dot1p</b>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>qos>network>egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command is used whenever the dot1p bits are set to a common value regardless of the internal profile of the packets. Although it is not mandatory, this command should be used in combination with the <b>de-mark</b> command to enable the marking of the DE bit according to the internal profile of the packet.</p> <p>This command acts as a shortcut for configuring the two existing commands with the same dot1p priority.</p> |

The **dot1p** *dot1p-priority* command is saved in the configuration as **dot1p-in-profile** *dot1p-priority* and **dot1p-out-profile** *dot1p-priority*. The inplus-profile traffic is marked with the same value as in-profile traffic. The exceed-profile traffic is marked with the same value as out-of-profile traffic.

## dot1p-in-profile

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dot1p-in-profile</b> <i>dot1p-priority</i><br><b>no dot1p-in-profile</b>                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>qos>network>egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command specifies dot1p in-profile mappings. The inplus-profile traffic is marked with the same value as in-profile traffic.</p> <p>The <b>no</b> form of this command resets the configuration to the default in-profile <i>dot1p-priority</i> setting for policy-id 1.</p>                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>dot1p-priority</i> — Specifies the unique IEEE 802.1p value that will match the dot1p rule. If the command is executed multiple times with the same <i>dot1p-value</i>, the previous forwarding class and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.</p> <p>A maximum of eight dot1p rules are allowed on a single policy.</p> <p><b>Values</b>      0 to 7</p> |

## dot1p-out-profile

|                    |                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dot1p-out-profile</b> <i>dot1p-priority</i><br><b>no dot1p-out-profile</b>                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>qos>network>egress>fc                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command specifies dot1p out-of-profile mappings.</p> <p>The exceed-profile traffic is marked with the same value as out-of-profile traffic.</p> <p>The <b>no</b> form of this command resets the configuration to the default out-profile <i>dot1p-priority</i> setting for policy-id 1.</p>                                                                                   |
| <b>Parameters</b>  | <p><i>dot1p-priority</i> — Specifies the unique IEEE 802.1p value that will match the dot1p rule. If the command is executed multiple times with the same <i>dot1p-value</i>, the previous forwarding class and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.</p> |

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

**Values**      0 to 7

dscp-in-profile

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dscp-in-profile</b> <i>dscp-name</i><br><b>no dscp-in-profile</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>network>egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command specifies the in-profile DSCP name for the forwarding class. The corresponding DSCP value will be used for all IP packets that require marking at egress on this forwarding class queue, and that are in-profile. The inplus-profile traffic is marked with the same value as in-profile traffic.</p> <p>When multiple DSCP names are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.</p> <p>The <b>no</b> form of this command resets the configuration to the factory default in-profile DSCP name setting for <i>policy-id</i> 1.</p> |
| <b>Parameters</b>  | <i>dscp-name</i> — Specifies the system- or user-defined, case-sensitive <i>dscp-name</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Values</b>      | Any defined system- or user-defined <i>dscp-name</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

dscp-out-profile

|                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |              |                 |                       |                   |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|-----------------|-----------------------|-------------------|
| <b>Syntax</b>         | <b>dscp-out-profile</b> <i>dscp-name</i><br><b>no dscp-out-profile</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |              |                 |                       |                   |
| <b>Context</b>        | config>qos>network>egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |              |                 |                       |                   |
| <b>Description</b>    | <p>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. The exceed-profile traffic is marked with the same value as out-of-profile traffic.</p> <p>When multiple DSCP names are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.</p> <p>The <b>no</b> form of this command resets the configuration to the factory default out-of-profile DSCP name setting for <i>policy-id</i> 1.</p> |              |                 |                       |                   |
| <b>Default</b>        | <table><tr><td>Policy-id 1:</td><td>Factory setting</td></tr><tr><td>Policy-id 2 to 65535:</td><td>Policy-id setting</td></tr></table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Policy-id 1: | Factory setting | Policy-id 2 to 65535: | Policy-id setting |
| Policy-id 1:          | Factory setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |              |                 |                       |                   |
| Policy-id 2 to 65535: | Policy-id setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |              |                 |                       |                   |



---

**Parameters**    *dscp-name* — Specifies the system- or user-defined, case-sensitive *dscp-name*.

**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>egress>fc

**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 that require marking at egress on this forwarding class queue, and that are in-profile. The inplus-profile traffic is marked with the same value as in-profile traffic.

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 resets the configuration to the factory default in-profile EXP setting.

**Default**

|                       |                   |
|-----------------------|-------------------|
| Policy-id 1:          | Factory setting   |
| Policy-id 2 to 65535: | Policy-id setting |

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

**Values**        0 to 7

## lsp-exp-out-profile

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

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

**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 that require marking at egress on this forwarding class queue, and that are out-of-profile. The exceed-profile traffic is marked with the same value as out-of-profile traffic.

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 resets the configuration to the factory default out-of-profile EXP setting.

---

|                   |                                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | Policy-id 1:                      Factory setting<br>Policy-id 2 to 65535:          Policy-id setting         |
| <b>Parameters</b> | <i>mpls-exp-value</i> — The 3-bit MPLS EXP bit value, expressed as a decimal integer.<br><b>Values</b> 0 to 7 |

## port-redirect-group

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>port-redirect-group</b> { <i>queue queue-id</i>   <i>policer policer-id</i> [ <i>queue queue-id</i> ]}<br><b>no port-redirect-group</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>network>egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command is used to redirect the FC of a packet of a pseudowire (PW) or network IP interface to an egress port queue group.</p> <p>It defines the mapping of an 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.</p> <p>The <b>no</b> version of this command removes the redirection of the FC.</p> |
| <b>Parameters</b>  | <p><i>queue-id</i> — This parameter must be specified when executing the <b>port-redirect-group</b> command. The specified <i>queue-id</i> must exist within the egress port queue group on each IP interface where the network QoS policy is applied.</p> <p>                  <b>Values</b>        1 to 8</p> <p><i>policer id</i> — The specified policer-id must exist within the queue-group template applied to the ingress context of the forwarding plane.</p> <p>                  <b>Values</b>        1 to 8</p>                                                                             |

## Network Ingress and Egress Criteria QoS Policy Commands

### ip-criteria

|                |                                                         |
|----------------|---------------------------------------------------------|
| <b>Syntax</b>  | <b>[no] ip-criteria</b>                                 |
| <b>Context</b> | config>qos>network>ingress<br>config>qos>network>egress |

**Description** IP criteria-based network ingress and egress policies are used to select the appropriate ingress or egress queue or policer, and the corresponding forwarding class and packet profile 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 DSCP.

The 7750 SR OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. Entries must be sequenced correctly from most to least explicit.

The ingress classification only applies to the outer IP header of non-tunneled traffic. The only exception is for traffic received on a Draft Rosen tunnel, for which only classification on the outer IP header is supported.

Attempting to apply a network QoS policy containing an **ip-criteria** statement to any object except a network IP interface will result in an error.

The **no** form of this command deletes all entries specified under this node. When IP criteria entries are removed from a network policy, the IP criteria are removed from all network interfaces to which that policy is applied.

## ipv6-criteria

**Syntax** **[no] ipv6-criteria**

**Context** config>qos>network>ingress  
config>qos>network>egress

**Description** IPv6 criteria-based network ingress and egress policies are used to select the appropriate ingress or egress queue or policer, and the corresponding forwarding class and packet profile for matched traffic. This command is used to enter the context to create or edit policy entries that specify IPv6 criteria such as IP quintuple lookup or DSCP.

The 7750 SR OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. Entries must be sequenced correctly from most to least explicit.

The ingress classification only applies to the outer IPv6 header of non-tunneled traffic.

Attempting to apply a network QoS policy containing an **ipv6-criteria** statement to any object except a network IP interface will result in an error.

The **no** form of this command deletes all entries specified under this node. When IP criteria entries are removed from a network policy, the IPv6 criteria are removed from all network interfaces to which that policy is applied.

## entry

**Syntax** **entry entry-id [create]**

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <b>no entry entry-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | <pre>config&gt;qos&gt;network&gt;ingress&gt;ip-criteria config&gt;qos&gt;network&gt;ingress&gt;ipv6-criteria config&gt;qos&gt;network&gt;egress&gt;ip-criteria config&gt;qos&gt;network&gt;egress&gt;ipv6-criteria</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>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 numbers.</p> <p>The list of flow criteria is evaluated in a top-down manner 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 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.</p> <p>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 ingress packets. If the <b>action</b> 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.</p> <p>The configuration of egress prec/DSCP classification and the configuration of an egress IP criteria or IPv6 criteria entry statement within a network QoS policy are mutually exclusive.</p> <p>Network QoS policies containing egress <b>ip-criteria</b> or <b>ipv6-criteria entry</b> statements are only applicable to network interfaces. Configuration of <b>ip-criteria</b> or <b>ipv6-criteria entry</b> statements in a network egress QoS policy and the application of the policy on any object other than a GRT network interface are mutually exclusive.</p> <p>The <b>no</b> form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services to which that policy is applied.</p> |
| <b>Parameters</b>  | <p><b>entry-id</b> — The entry identifier, expressed as an integer, uniquely identifies a match criterion and the corresponding action. It is recommended that multiple entries be given entry identifiers in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.</p> <p>An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword <b>action fc fc-name profile profile</b> for it to be considered complete. Entries without the action keyword will be considered incomplete and will be rendered inactive.</p> <p><b>Values</b>      1 to 65535</p> <p><b>create</b> — 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.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## action

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>action</b> [ <b>fc</b> <i>fc-name</i> <b>profile</b> { <b>in</b>   <b>out</b>   <b>exceed</b>   <b>inplus</b> }] [ <b>port-redirect-group</b> { <b>queue</b> <i>queue-id</i>   <b>policer</b> <i>policer-id</i> [ <b>queue</b> <i>queue-id</i> ]}]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>qos>network>egress>ip-criteria>entry<br>config>qos>network>egress>ipv6-criteria>entry                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command defines the reclassification actions that are performed on any packet matching the defined IP flow criteria within the entry's matched node. When defined under the <b>ip-criteria</b> context, the reclassification only applies to IPv4 packets. When defined under the <b>ipv6-criteria</b> context, the reclassification only applies to IPv6 packets.</p> <p>If an egress packet matches the specified IP flow entry, the forwarding class and profile may be overridden. By default, the forwarding class and profile of the packet are derived from ingress classification and profiling functions. Matching an IP flow reclassification entry will override all IP precedence-based or DSCP-based reclassification rule actions when an explicit reclassification action is defined for the entry.</p> <p>When an IP flow entry is first created, the entry will have no explicit behavior defined as the reclassification actions to be performed. When the entry is defined with no action, the entry will not be populated in the IP flow reclassification list used to evaluate egress packets. An IP flow reclassification entry is only added to the evaluation list when the action command for the entry is executed.</p> <p>The <b>fc</b> and <b>profile</b> keywords are optional. When specified, the egress classification rule will overwrite the forwarding class and profile derived from ingress. The new forwarding class and profile are used for egress remarking, queue mapping decisions, and queue congestion behavior.</p> <p>The <b>port-redirect-group</b> keyword is optional. When specified, the egress packet will be redirected to the configured queue or policer in the specified egress network queue group. By default, the policed forwarded traffic will use the regular network queue to which the packet's forwarding class is mapped. Alternatively, a queue in the network egress queue group instance can be used for post-policed traffic by specifying a queue after the <b>policer</b> parameter. The <b>port-redirect-group</b> keyword requires that the network egress queue group instance is specified when this network QoS policy is applied to a network interface. The <b>port-redirect-group</b> is not supported on a 7750 SR-a4/a8.</p> <p>The <b>no</b> 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 egress packets.</p> |
| <b>Default</b>     | no action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**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. The profile reclassification action is mandatory when an **fc** is specified. 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**        be, l2, af, l1, h2, ef, h1, nc

**profile {in | out | exceed | inplus}** — The profile reclassification action is mandatory when an fc is specified, otherwise it is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to the configured profile regardless of the ingress profiling decision. **In**, **exceed**, **inplus**, or **out** must be specified when the profile keyword is present. 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** — When specified, any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

**out** — When specified, any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

**exceed** — When specified, any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

**inplus** — When specified, any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

**queue** *queue-id* — Used to override the forwarding class default egress queue destination to the specified network egress queue group instance queue. The specific egress queue group instance to use is specified at the time the QoS policy is applied to the network interface.

**Values**        1 to 8

**policer** *policer-id* — When policer parameter is entered, a valid policer ID must be specified. The parameter policer ID references a policer ID that has already been created within the network egress queue group instance.

**Values**        1 to 8

**queue** *queue-id* — The queue following the configured policer overrides the default policed traffic egress queue destination to a specified queue in the network egress queue group instance.

**Values**        1 to 8

## action

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

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <b>no action</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>qos>network>ingress>ip-criteria>entry<br>config>qos>network>ingress>ipv6-criteria>entry                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command defines the reclassification actions that are performed on any packet matching the defined IP flow criteria within the entry's matched node. When defined under the <b>ip-criteria</b> context, the reclassification only applies to IPv4 packets. When defined under the <b>ipv6-criteria</b> context, the reclassification only applies to IPv6 packets.</p> <p>If an ingress packet matches the specified IP flow entry, the forwarding class and profile may be overridden. By default, the forwarding class and profile of the packet are derived from ingress classification and profiling functions. Matching an IP flow reclassification entry will override all non-criteria reclassification rule actions when an explicit reclassification action is defined for the entry.</p> <p>When an IP flow entry is first created, the entry will have no explicit behavior defined as the reclassification actions to be performed. When the entry is defined with no action, the entry will not be populated in the IP flow reclassification list used to evaluate ingress packets. An IP flow reclassification entry is only added to the evaluation list when the action command for the entry is executed.</p> <p>The <b>no</b> 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 ingress packets.</p>                            |
| <b>Default</b>     | no action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <p><b>fc</b> <i>fc-name</i> — 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 <i>fc-name</i> regardless of the ingress classification decision. The <i>fc-name</i> defined must be one of the eight forwarding classes supported by the system. The profile reclassification action is mandatory when an <b>fc</b> is specified. 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.</p> <p><b>Values</b>      be, l2, af, l1, h2, ef, h1, nc</p> <p><b>profile {in   out}</b> — The profile reclassification action is mandatory. Packets matching the IP flow reclassification entry will be explicitly reclassified to the configured profile regardless of other ingress profiling decisions. <b>In</b> or <b>out</b> must be specified when the profile keyword is present. 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.</p> <p><b>in</b> — When specified, any packets matching the reclassification rule will be treated as in-profile by the ingress forwarding plane.</p> <p><b>out</b> — When specified, any packets matching the reclassification rule will be treated as out-of-profile by the ingress forwarding plane.</p> |

## match

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>match</b> [ <b>protocol</b> <i>protocol-id</i> ]<br><b>no match</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>qos>network>ingress>ip-criteria>entry<br>config>qos>network>egress>ip-criteria>entry                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command creates a context to configure match criteria for a network QoS policy. When the match criteria have been satisfied, the action associated with it is executed.</p> <p>If more than one match criteria (within one match statement) are configured, then all criteria must be satisfied before the associated action with the match is executed.</p> <p>A match context can consist of multiple match criteria, but multiple match statements cannot be entered per entry.</p> <p>A network QoS policy can include the DSCP map command, the dot1p map command (ingress only), the prec map command (egress only), 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:</p> <ul style="list-style-type: none"> <li>• 802.1p bits (ingress only)</li> <li>• DSCP</li> <li>• prec (egress only)</li> <li>• IP quintuple</li> </ul> <p>The <b>no</b> form of this command removes the match criteria for the entry identifier.</p> |
| <b>Parameters</b>  | <p><b>protocol</b> <i>protocol-id</i> — Specifies an IP protocol to be used as an ingress or egress network QoS policy match criterion.</p> <p>The protocol type is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), and UDP(17).</p> <p><b>Values</b>     <i>protocol-id</i>: 0 to 255 protocol numbers accepted in decimal, hexadecimal, or binary</p> <p>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</p> <p>* — udp/tcp wildcard</p> <p><a href="#">Table 21</a> lists the protocols and their protocol IDs and descriptions.</p>                                                                                                                                                                                                                                                                                       |

**Table 21**      Protocol ID Descriptions

| Protocol | Protocol ID | Description              |
|----------|-------------|--------------------------|
| icmp     | 1           | Internet Control Message |



**Table 21 Protocol ID Descriptions (Continued)**

| Protocol    | Protocol ID | Description                                                 |
|-------------|-------------|-------------------------------------------------------------|
| 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                                |
| stp         | 118         | Schedule Transfer Protocol                                  |
| ptp         | 123         | Performance Transparency Protocol                           |

**Table 21** Protocol ID Descriptions (Continued)

| Protocol | Protocol ID | Description                     |
|----------|-------------|---------------------------------|
| isis     | 124         | ISIS over IPv4                  |
| crtip    | 126         | Combat Radio Transport Protocol |
| crudp    | 127         | Combat Radio User Datagram      |

## match

**Syntax** **match** [*next-header next-header*]  
**no match**

**Context** config>qos>network>ingress>ipv6-criteria>entry  
config>qos>network>egress>ipv6-criteria>entry

**Description** This command creates a context to configure match criteria for a network 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, 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.

A network policy can include the DSCP map command, the dot1p map command (ingress only), the prec map command (egress only), 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:

- 802.1p bits (ingress only)
- DSCP
- prec (egress only)
- IP quintuple

The **no** form of this command removes the match criteria for the entry identifier.

**Parameters** **next-header next-header** — Specifies the next header to match.

The protocol type is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), and UDP(17).

**Values** protocol numbers accepted in decimal, hexadecimal, or binary: 0 to 42, 45 to 49, 52 to 59, 61 to 255  
keywords: none, crt, crudp, ecp, 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

## dscp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dscp</b> <i>dscp-name</i><br><b>no dscp</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>qos>network>ingress>ip-criteria>entry>match<br>config>qos>network>ingress>ipv6-criteria>entry>match<br>config>qos>network>egress>ip-criteria>entry>match<br>config>qos>network>egress>ipv6-criteria>entry>match                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command configures a DSCP to be used as a network QoS policy match criterion.<br><br>The <b>no</b> form of this command removes the DSCP match criterion.                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <i>dscp-name</i> — Specifies a DSCP name that has been previously mapped to a value using the <b>dscp-name</b> command. The DSCP can only be specified by its name.<br><br><b>Values</b> 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 |

## dst-ip

|                    |                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dst-ip</b> { <i>ip-address/mask</i>   <i>ip-address ipv4-address-mask</i>   <b>ip-prefix-list</b> <i>ip-prefix-list-name</i> }<br><b>dst-ip</b> { <i>ipv6-address/mask</i>   <i>ipv6-address ipv6-address-mask</i>   <b>ipv6-prefix-list</b> <i>ipv6-prefix-list-name</i> }<br><b>no dst-ip</b>                                                                                            |
| <b>Context</b>     | config>qos>network>ingress>ip-criteria>entry>match<br>config>qos>network>ingress>ipv6-criteria>entry>match                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | This command configures a destination address range to be used as a network QoS policy match criterion.<br><br>To match on the destination address, specify the address and its associated mask, for example, when specifying an IPv4 address, 10.1.0.0/16 or 10.1.0.0 255.255.0.0 can be used.<br><br>The <b>no</b> form of this command removes the destination IP address match criterion. |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the source IPv4 address specified in dotted decimal notation.<br><br><b>Values</b> ip-address: a.b.c.d                                                                                                                                                                                                                                                          |

*mask* — Specifies the length in bits of the subnet mask.

**Values** 1 to 32

*ipv4-address-mask* — Specifies the subnet mask in dotted decimal notation.

**Values** a.b.c.d (dotted quad equivalent of mask length)

*ip-prefix-list-name* — Specifies an IPv4 prefix list which contains IPv4 address prefixes to be matched. IP prefix lists are only supported at a network ingress.

**Values** A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

*ipv6-address* — Specifies the IPv6 prefix for the IP match criterion in hex digits.

**Values**

ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces)  
x:x:x:x:x:x:d.d.d.d  
x: [0 to FFFF]H  
d: [0 to 255]D

*mask* — Specifies the length of the IPv6 address expressed as a decimal integer.

**Values** 1 to 128

*ipv6-address-mask* — Specifies the eight 16-bit hexadecimal pieces representing bit match criteria.

**Values** x:x:x:x:x:x (eight 16-bit pieces)

*ipv6-prefix-list-name* — Specifies an IPv6 prefix list which contains IPv6 address prefixes to be matched. IPv6 prefix lists are only supported at a network ingress.

**Values** A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

## dst-port

|                    |                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dst-port</b> {lt   gt   eq} <i>dst-port-number</i><br><b>dst-port range</b> <i>start end</i><br><b>no dst-port</b>                                                                                                  |
| <b>Context</b>     | config>qos>network>ingress>ip-criteria>entry>match<br>config>qos>network>ingress>ipv6-criteria>entry>match<br>config>qos>network>egress>ip-criteria>entry>match<br>config>qos>network>egress>ipv6-criteria>entry>match |
| <b>Description</b> | This command configures a destination TCP or UDP port number or port range for a network QoS policy match criterion.                                                                                                   |

The **no** form of this command removes the destination port match criterion.

- 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) the destination port value, specified as a decimal integer.
- Values** 1 to 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 to 65535 (decimal)

## fragment

- Syntax** **fragment {true | false}**  
**no fragment**
- Context** config>qos>network>ingress>ip-criteria>entry>match  
config>qos>network>egress>ip-criteria>entry>match
- Description** This command configures fragmented or non-fragmented IP packets as a network 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.
- 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.

## fragment

- Syntax** **fragment {true | false | first-only | non-first-only}**  
**no fragment**
- Context** config>qos>network>ingress>ipv6-criteria>entry>match  
config>qos>network>egress>ipv6-criteria>entry>match
- Description** This command configures fragmented or non-fragmented IPv6 packets as a network 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.

- Parameters**
- true** — Specifies to match on all fragmented IPv6 packets. A match will occur for all packets that contain an IPv6 Fragmentation Extension Header.
  - false** — Specifies to match on all non-fragmented IP packets. Non-fragmented IPv6 packets are packets that do not contain an IPv6 Fragmentation Extension Header.
  - first-only** — Matches if a packet is an initial fragment of the fragmented IPv6 packet.
  - non-first-only** — Matches if a packet is a non-initial fragment of the fragmented IPv6 packet.

## icmp-type

- Syntax** **icmp-type** *icmp-type*  
**no icmp-type**
- Context** config>qos>network>egress>ip-criteria>entry>match  
config>qos>network>egress>ipv6-criteria>entry>match
- Description** This command configures matching on the ICMP or ICMPv6 type field in the ICMP or ICMPv6 header of an IPv4 or IPv6 packet as a network QoS match criterion.
- An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, etc.) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. Similarly, an entry containing "**icmp-type 0**" match criterion, may match non-initial fragments when the Layer 4 header is not present in a packet fragment and other match criteria are also met.
- The **no** form of the command removes the criterion from the match entry.
- Default** no icmp-type
- Parameters** *icmp-type* — Specifies the ICMP or ICMPv6 type value that must be present to match. Value can be expressed as a decimal integer, or in hexadecimal or binary format, or even using keywords.
- Values**
- 0 to 255 (Decimal)
  - 0 to FF (Hexadecimal)
  - 0 to 11111111 (Binary)

## src-ip

- Syntax** **src-ip** {*ip-address/mask* | *ip-address ipv4-address-mask* | **ip-prefix-list** *ip-prefix-list-name*}  
**src-ip** {*ipv6-address/mask* | *ipv6-address ipv6-address-mask* | **ipv6-prefix-list** *ipv6-prefix-list-name*}  
**no src-ip**
- Context** config>qos>network>ingress>ip-criteria>entry>match  
config>qos>network>ingress>ipv6-criteria>entry>match

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <pre>config&gt;qos&gt;network&gt;egress&gt;ip-criteria&gt;entry&gt;match config&gt;qos&gt;network&gt;egress&gt;ipv6-criteria&gt;entry&gt;match</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command configures a source IPv4 or IPv6 address range to be used as a network QoS policy match criterion.</p> <p>To match on the source IPv4 or IPv6 address, specify the address and its associated mask, for example, when specifying an IPv4 address, 10.1.0.0/16 or 10.1.0.0 255.255.0.0 can be used.</p> <p>The <b>no</b> form of this command removes the source IPv4 or IPv6 address match criterion.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>ip-address</i> — Specifies the source IPv4 address specified in dotted decimal notation.</p> <p><b>Values</b> ip-address: a.b.c.d</p> <p><i>mask</i> — Specifies the length in bits of the subnet mask.</p> <p><b>Values</b> 1 to 32</p> <p><i>ipv4-address-mask</i> — Specifies the subnet mask in dotted decimal notation.</p> <p><b>Values</b> a.b.c.d (dotted quad equivalent of mask length)</p> <p><i>ip-prefix-list-name</i> — Specifies an IPv4 prefix list which contains IPv4 address prefixes to be matched. IP prefix lists are only supported at a network ingress.</p> <p><b>Values</b> A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.</p> <p><i>ipv6-address</i> — Specifies the IPv6 prefix for the IP match criterion in hex digits.</p> <p><b>Values</b></p> <pre>ipv6-address:  x:x:x:x:x:x:x (eight 16-bit pieces)                 x:x:x:x:x:d.d.d.d                 x: [0 to FFFF]H                 d: [0 to 255]D</pre> <p><i>mask</i> — Specifies the length of the ipv6-address expressed as a decimal integer.</p> <p><b>Values</b> 1 to 128</p> <p><i>ipv6-address-mask</i> — Specifies the eight 16-bit hexadecimal pieces representing bit match criteria.</p> <p><b>Values</b> x:x:x:x:x:x (eight 16-bit pieces)</p> <p><i>ipv6-prefix-list-name</i> — Specifies an IPv6 prefix list which contains IPv6 address prefixes to be matched. IPv6 prefix lists are only supported at a network ingress</p> <p><b>Values</b> A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.</p> |

---

## src-port

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>src-port</b> {lt   gt   eq} <i>src-port-number</i><br><b>src-port range</b> <i>start end</i><br><b>no src-port</b>                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>network>ingress>ip-criteria>entry>match<br>config>qos>network>ingress>ipv6-criteria>entry>match<br>config>qos>network>egress>ip-criteria>entry>match<br>config>qos>network>egress>ipv6-criteria>entry>match                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command configures a source TCP or UDP port number or port range for a network QoS policy match criterion.</p> <p>The <b>no</b> form of this command removes the source port match criterion.</p>                                                                                                                                                                                                                                                          |
| <b>Default</b>     | No src-port match criterion.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <p><b>lt   gt   eq</b> <i>src-port-number</i> — 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.</p> <p><b>Values</b> 1 to 65535 (decimal)</p> <p><b>range</b> <i>start end</i> — The range of TCP or UDP port values to match, specified as between the start and end source port values inclusive.</p> <p><b>Values</b> 1 to 65535 (decimal)</p> |

## renum

|                    |                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>renum</b> <i>old-entry-number new-entry-number</i>                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>network>ingress>ip-criteria<br>config>qos>network>ingress>ipv6-criteria<br>config>qos>network>egress>ip-criteria<br>config>qos>network>egress>ipv6-criteria                                                                                                                                                                                             |
| <b>Description</b> | <p>This command renumbers existing QoS policy criteria entries to properly sequence policy entries.</p> <p>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.</p> |
| <b>Parameters</b>  | <p><b>old-entry-number</b> — Enter the entry number of an existing entry.</p> <p><b>Values</b> 1 to 65535</p> <p><b>new-entry-number</b> — Enter the new entry number to be assigned to the old entry.</p> <p><b>Values</b> 1 to 65535</p>                                                                                                                         |



#### 4.7.2.1.4 Model-Driven Automatic ID Commands

##### md-auto-id

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>md-auto-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>qos                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command automatically assigns numerical ID values for model-driven (MD) management interfaces.</p> <p>Classic management interfaces use a numerical policy ID as the primary key for sap-ingress, sap-egress, and network QoS policies. In model-driven interfaces, SAP and network policies use string names as keys. The SAP and network policies can optionally be created without having to explicitly select and specify a numerical policy ID. In this case, SR OS assigns an ID using the configured ID range.</p> |

##### qos-policy-id-range

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>qos-policy-id-range start <i>policy-id</i> end <i>policy-id</i></b><br><b>no qos-policy-id-range</b>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>qos>md-auto-id                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command specifies the range of IDs used by SR OS to automatically assign an ID to QoS policies that are created in model-driven interfaces without an ID explicitly specified by the user or client.</p> <p>A QoS policy created with an explicitly-specified ID cannot use an ID in this range. The ID range cannot be changed while policies exist inside the previous or new range.</p> <p>The <b>no</b> form of this command removes the range values.</p> <p>See the <a href="#">md-auto-id</a> command for further details.</p> |
| <b>Default</b>     | no qos-policy-id-range                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><b>start <i>policy-id</i></b> — Specifies the lower value of the ID range. The value must be less than or equal to the <b>end</b> value.</p> <p><b>Values</b> 1 to 65535</p> <p><b>end <i>policy-id</i></b> — Specifies the upper value of the ID range. The value must be greater than or equal to the <b>start</b> value.</p> <p><b>Values</b> 1 to 65535</p>                                                                                                                                                                            |

---

### 4.7.2.2 Self-Generated Traffic Commands

#### sgt-qos

|                    |                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sgt-qos</b>                                                                                |
| <b>Context</b>     | config>router<br>config>service>vprn                                                          |
| <b>Description</b> | This command enters the context to configure DSCP/dot1p remarking for self-generated traffic. |

#### application

|                    |                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>application</b> <i>dscp-app-name</i> <b>dscp</b> { <i>dscp-value</i>   <i>dscp-name</i> }<br><b>application</b> <i>dot1p-app-name</i> <b>dot1p</b> <i>dot1p-priority</i><br><b>no application</b> { <i>dscp-app-name</i>   <i>dot1p-app-name</i> }                                                                              |
| <b>Context</b>     | config>router>sgt-qos<br>config>service>vprn>sgt-qos                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command configures DSCP/dot1p remarking for self-generated application traffic. When an application is configured using this command, 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 Ethernet 802.1p and MPLS EXP bits.
- The Ethernet 802.1p and MPLS 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, the subsequent entry overrides the previous configured entry.

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

|                   |                                                             |
|-------------------|-------------------------------------------------------------|
| <b>Parameters</b> | <i>dscp-app-name</i> — Specifies the DSCP application name. |
| <b>Values</b>     | The following values apply to the 7750 SR and 7950 XRS:     |

bgp, call-trace, cflowd, dhcp, diameter, dns, ftp, gtp, icmp, igmp, igmp-reporter, l2tp, ldp, mld, msdp, ndis, ntp, ospf, pcep, pim, ptp, radius, rip, rsvp, sflow, snmp, snmp-notification, srtp, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp

**Values** The following values apply to the 7450 ESS:

bgp, call-trace, dhcp, dns, ftp, icmp, igmp, ldp, mld, ntp, ospf, radius, rsvp, rip, ssh, snmp, snmp-notification, syslog, telnet, tftp, traceroute, tacplus, vrrp

*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 to 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 to 7

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

**Values** The following values apply to the 7750 SR and 7950 XRS:  
arp, isis, pppoe

**Values** The following values apply to the 7450 ESS  
arp, isis

## dscp

|                    |                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dscp</b> <i>dscp-name</i> <b>fc</b> <i>fc-name</i><br><b>no dscp</b> <i>dscp-name</i>                                      |
| <b>Context</b>     | config>router>sgt-qos<br>config>service>vpn>sgt-qos                                                                           |
| <b>Description</b> | 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 64 DSCPs to the forwarding class.

All DSCP names that define a DSCP value must be explicitly defined

The **no** form of this command removes the DSCP-to-forwarding class association.

**Parameters** *dscp-name* — Specifies the name of the DSCP 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* — Specifies the forwarding class name. All packets with a DSCP value or MPLS EXP bit that are not defined will be placed in this forwarding class.

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

### 4.7.2.3 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 to 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
```

## 4.7.2.4 Show Commands

### dscp-table

|                    |                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dscp-table</b> [ <b>value</b> <i>dscp-value</i> ]                                                                                                  |
| <b>Context</b>     | show>qos                                                                                                                                              |
| <b>Description</b> | This command displays the DSCP name-to-DSCP value mappings.                                                                                           |
| <b>Parameters</b>  | <b>value</b> <i>dscp-value</i> — The specific DSCP value for which to display information.<br><b>Values</b> 0 to 63<br><b>Default</b> Show all values |
| <b>Output</b>      | The following output is an example of dscp-table information, and <a href="#">Table 22</a> describes the output fields.                               |

#### 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#

```

**Table 22** Show QoS DSCP Table Output Fields

| Label     | Description                                                                              |
|-----------|------------------------------------------------------------------------------------------|
| DSCP Name | Displays the name of the DiffServ code point to be associated with the forwarding class. |

**Table 22 Show QoS DSCP Table Output Fields (Continued)**

| Label      | Description                                    |
|------------|------------------------------------------------|
| DSCP Value | Displays the DSCP values, between 0 and 63.    |
| TOS (bin)  | Displays the type of service in binary format. |
| TOS (hex)  | Displays the type of service in hex format.    |

## mc-fr-profile-ingress

- Syntax** **mc-fr-profile-ingress**  
**mc-fr-profile-ingress** *profile-id* [**detail**]
- Context** show>qos
- Description** This command displays MLFR ingress profile information.
- Output** The following output is an example of MLFR ingress profile information.

### 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**  
**mc-fr-profile-egress** *profile-id* [**detail**]
- Context** show>qos
- Description** This command displays MLFR egress profile details.
- Output** The following output is an example of MLFR ingress profile information.

### 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#
```

## md-auto-id

- Syntax** **md-auto-id**
- Context** show>qos
- Description** This command displays MD automatically-assigned ID information.
- Output** The following output is an example of MD automatically-assigned QoS information, and [Table 23](#) describes the MD automatic ID QoS output fields.



### Sample Output

```
*A:node-6# show qos md-auto-id
=====
MD Auto-Id Information
=====
QoS Policy-Id Range
Start : 327675 End : 65535
Sap Ingress Count : 10
Sap Egress Count : 8
Network Count : 6
=====
```

**Table 23 MD Auto-Id QoS Output Fields**

| Label             | Description                                                                   |
|-------------------|-------------------------------------------------------------------------------|
| Start             | Specifies the start range for a QoS policy ID.                                |
| End               | Specifies the end range for a QoS Policy ID.                                  |
| Sap Ingress Count | Specifies the number of SAP ingress policies with automatically-assigned IDs. |
| Sap Egress Count  | Specifies the number of SAP egress policies with automatically-assigned IDs.  |
| Network Count     | Specifies the number of network policies with automatically-assigned IDs.     |

## network

- Syntax** **network** [*policy-id*] [**detail**] [**match-criteria**]
- Context** show>qos
- Description** This command displays network policy information.
- Parameters** *policy-id* — Displays information for the specific policy ID.
- Values** 1 to 65535
- Default** all network policies
- detail** — Includes information about ingress and egress DSCP and LSP EXP bit mappings and network policy interface associations.
- match-criteria** — Includes the match-criteria.
- Output** The following output is an example of network QoS Policy information, and [Table 24](#) describes the network QoS Policy output fields.

**Sample Output**

```

A:PE1# show qos network
=====
Network Policies
=====
Policy-Id Name Remark LerUseDscp Description

1 default False False Default network QoS policy.
=====
A:PE1#

A:PE1# show qos network 1
=====
QoS Network Policy
=====
Network Policy (1)

Policy-id : 1 Remark : False
Forward Class : be Profile : Out
LER Use DSCP : False Scope : Template
Description : Default network QoS policy.
Name : default
=====
A:PE1#

A:PE1# show qos network 1 detail
=====
QoS Network Policy
=====
Network Policy (1)

Policy-id : 1 Remark : False
Forward Class : be Profile : Out
LER Use DSCP : False Scope : Template
Description : Default network QoS policy.
Name : default
=====
DSCP (Ingress) Forwarding Class Profile

be be Out
ef ef In
cs1 l2 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

DSCP (Egress) Forwarding Class Profile

No Matching Entries

Prec (Egress) Forwarding Class Profile

No Matching Entries

LSP EXP Bit Map Forwarding 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

Dot1p Bit Map Forwarding Class Profile

No Matching Entries

Egress Forwarding Class Mapping

FC Value : 0 FC Name : be
- DSCP Mapping
Out-of-Profile : be In-Profile : be

- Dot1p Mapping
Out-of-Profile : 0 In-Profile : 0

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

DE Mark : None
Redirect Grp Q : None Redirect Grp Plcr: None

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
Redirect Grp Q : None Redirect Grp Plcr: None

FC Value : 2 FC Name : af

```

---

|                       |        |                    |        |
|-----------------------|--------|--------------------|--------|
| - DSCP Mapping        |        |                    |        |
| Out-of-Profile        | : af12 | In-Profile         | : af11 |
| - Dot1p Mapping       |        |                    |        |
| Out-of-Profile        | : 2    | In-Profile         | : 2    |
| - LSP EXP Bit Mapping |        |                    |        |
| Out-of-Profile        | : 2    | In-Profile         | : 3    |
| DE Mark               | : None |                    |        |
| Redirect Grp Q        | : None | Redirect Grp Plcr: | None   |
| FC Value              | : 3    | FC Name            | : l1   |
| - DSCP Mapping        |        |                    |        |
| Out-of-Profile        | : af22 | In-Profile         | : af21 |
| - Dot1p Mapping       |        |                    |        |
| Out-of-Profile        | : 3    | In-Profile         | : 3    |
| - LSP EXP Bit Mapping |        |                    |        |
| Out-of-Profile        | : 2    | In-Profile         | : 3    |
| DE Mark               | : None |                    |        |
| Redirect Grp Q        | : None | Redirect Grp Plcr: | None   |
| FC Value              | : 4    | FC Name            | : h2   |
| - DSCP Mapping        |        |                    |        |
| Out-of-Profile        | : af42 | In-Profile         | : af41 |
| - Dot1p Mapping       |        |                    |        |
| Out-of-Profile        | : 4    | In-Profile         | : 4    |
| - LSP EXP Bit Mapping |        |                    |        |
| Out-of-Profile        | : 4    | In-Profile         | : 4    |
| DE Mark               | : None |                    |        |
| Redirect Grp Q        | : None | Redirect Grp Plcr: | None   |
| FC Value              | : 5    | FC Name            | : ef   |
| - DSCP Mapping        |        |                    |        |
| Out-of-Profile        | : ef   | In-Profile         | : ef   |
| - Dot1p Mapping       |        |                    |        |
| Out-of-Profile        | : 5    | In-Profile         | : 5    |
| - LSP EXP Bit Mapping |        |                    |        |
| Out-of-Profile        | : 5    | In-Profile         | : 5    |
| DE Mark               | : None |                    |        |
| Redirect Grp Q        | : None | Redirect Grp Plcr: | None   |
| FC Value              | : 6    | FC Name            | : h1   |
| - DSCP Mapping        |        |                    |        |
| Out-of-Profile        | : nc1  | In-Profile         | : nc1  |
| - Dot1p Mapping       |        |                    |        |
| Out-of-Profile        | : 6    | In-Profile         | : 6    |
| - LSP EXP Bit Mapping |        |                    |        |

|                       |                         |
|-----------------------|-------------------------|
| Out-of-Profile : 6    | In-Profile : 6          |
| DE Mark : None        |                         |
| Redirect Grp Q : None | Redirect Grp Plcr: None |
| FC Value : 7          | FC Name : nc            |
| - DSCP Mapping        |                         |
| Out-of-Profile : nc2  | In-Profile : nc2        |
| - Dot1p Mapping       |                         |
| Out-of-Profile : 7    | In-Profile : 7          |
| - LSP EXP Bit Mapping |                         |
| Out-of-Profile : 7    | In-Profile : 7          |
| DE Mark : None        |                         |
| Redirect Grp Q : None | Redirect Grp Plcr: None |

-----

-----  
Ingress Forwarding Class Mapping

|                                |                                |
|--------------------------------|--------------------------------|
| FC Value : 0                   | FC Name : be                   |
| Redirect UniCast Plcr : None   | Redirect MultiCast Plcr : None |
| Redirect BroadCast Plcr : None | Redirect Unknown Plcr : None   |
| FC Value : 1                   | FC Name : l2                   |
| Redirect UniCast Plcr : None   | Redirect MultiCast Plcr : None |
| Redirect BroadCast Plcr : None | Redirect Unknown Plcr : None   |
| FC Value : 2                   | FC Name : af                   |
| Redirect UniCast Plcr : None   | Redirect MultiCast Plcr : None |
| Redirect BroadCast Plcr : None | Redirect Unknown Plcr : None   |
| FC Value : 3                   | FC Name : l1                   |
| Redirect UniCast Plcr : None   | Redirect MultiCast Plcr : None |
| Redirect BroadCast Plcr : None | Redirect Unknown Plcr : None   |
| FC Value : 4                   | FC Name : h2                   |
| Redirect UniCast Plcr : None   | Redirect MultiCast Plcr : None |
| Redirect BroadCast Plcr : None | Redirect Unknown Plcr : None   |
| FC Value : 5                   | FC Name : ef                   |
| Redirect UniCast Plcr : None   | Redirect MultiCast Plcr : None |
| Redirect BroadCast Plcr : None | Redirect Unknown Plcr : None   |
| FC Value : 6                   | FC Name : h1                   |
| Redirect UniCast Plcr : None   | Redirect MultiCast Plcr : None |
| Redirect BroadCast Plcr : None | Redirect Unknown Plcr : None   |
| FC Value : 7                   | FC Name : nc                   |
| Redirect UniCast Plcr : None   | Redirect MultiCast Plcr : None |
| Redirect BroadCast Plcr : None | Redirect Unknown Plcr : None   |

-----

-----  
Match Criteria (Ingress)

```

No Matching Entries

```

```

Match Criteria (Egress)

```

```

No Matching Entries

```

```

Interface Association

```

```
Interface : system
IP Addr. : n/a
Interface : int
IP Addr. : 20.0.0.1/24
Port Id : system
Port Id : 1/1/1
```

```

Service Associations

```

```
Service Id Service Type Direction

No Matching Entries Found

```

```

SDP Associations

```

```
Service Id Service Type Sdp Id Direction

No Matching Entries Found

```

```
=====
*A:PE1#
```

**Table 24** 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.<br>False — Remarking is disabled. |
| Description | A text string that helps identify the policy's context in the configuration file.                                                                                                                                                                                                  |

**Table 24 Show QoS Network Output Fields (Continued)**

| Label                        | Description                                                                                                                                                                                                     |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Class/<br>FC Name | Specifies the default ingress forwarding class (configured using the <b>network ingress default-action</b> command) assigned to packets for which there is no other configured forwarding class classification. |
| Profile                      | Specifies the default ingress packet profile (configured using the <b>network ingress default-action</b> command) assigned to packets for which there is no other configured packet profile classification.     |
| 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.                                                                                                                                                                                    |
| IP Addr                      | Displays the interface IP address.                                                                                                                                                                              |
| Port-Id                      | Specifies the physical port identifier that associates the interface.                                                                                                                                           |

## sgt-qos

|                    |                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sgt-qos</b>                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | show>router                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command displays QoS information for self-generated traffic. In the output, “none” means that the default values for each application are used, not that there is no value set. For a list of application defaults, see <a href="#">QoS for Self-Generated (CPU) Traffic on Network Interfaces</a> and <a href="#">Table 16</a> . |

## application

|                    |                                                 |
|--------------------|-------------------------------------------------|
| <b>Syntax</b>      | <b>application</b> [app-name] [dscp   dot1p]    |
| <b>Context</b>     | show>router>sgt-qos                             |
| <b>Description</b> | This command displays application QoS settings. |

---

|                   |                                                                                                                                                                                                                                                                                                                          |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>app-name</i> — The specific application.                                                                                                                                                                                                                                                                              |
| <b>Values</b>     | The following values apply to the 7750 SR and 7950 XRS:<br>arp, bgp, cflowd, dhcp, diameter, dns, ftp, gtp, icmp, igmp, igmp-reporter, isis, l2tp, ldp, mld, msdp, ndis, ntp, ospf, pcep, pim, pppoe, ptp, radius, rip, rsvp, sflow, snmp, snmp-notification, srrp, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp |
| <b>Values</b>     | The following values apply to the 7450 ESS:<br>arp, bgp, dhcp, dns, ftp, icmp, isis, ldp, ntp, ospf, radius, rip, rsvp, snmp, snmp-notification, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp                                                                                                                    |

## dscp-map

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dscp-map</b> [ <i>dscp-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | show>router>sgt-qos                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | This command displays DSCP to FC mappings.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <i>dscp-name</i> — The specific DSCP name.<br>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 |



## 5 Network Queue QoS Policies

### 5.1 Overview

Network queue policies define the ingress network queuing at the XMA/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 XMA/MDA network ingress ports. Other network queue QoS policies must be explicitly created and associated.

For information about the tasks and commands necessary to access the CLI and to configure and maintain devices, refer to the CLI Usage chapter in the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Basic System Configuration Guide*.

---

## 5.2 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 that exists on a port without a scheduler policy defined will be considered to be orphaned.

Refer to [Scheduler QoS Policies](#) for more information about queue parental association scope.

---

## 5.3 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.

### 5.3.1 Creating 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 XMA and MDA network ingress ports.

To create a network queue policy, enter the **configure qos network-queue** command with an unused policy name followed by the **create** keyword. Once created, the parameters in the policy, such as a description, the mapping of forwarding classes to queues, and the queue properties can be configured.

### 5.3.2 Applying Network Queue QoS Policies

Apply network queue policies to the following entities:

- XMAs and MDAs
- Ethernet Ports
- SONET/SDH Ports

#### 5.3.2.1 XMAs and MDAs

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

**CLI Syntax:**

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

The following output displays XMA or 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
access
 ingress
 pool default
 resv-cbs sum
 slope-policy "default"
 exit
 egress
 pool default
 resv-cbs sum
 slope-policy "default"
 exit
exit
exit
no shutdown

A:ALA-7>config>card>mda#
```

### 5.3.2.2 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
```

The following output displays the port configuration.

```
A:ALA-49>config>port# info

ethernet
 network
 queue-policy "nq1"
```

```

 exit
 exit
 no shutdown

A:ALA-49>config>port#

```

### 5.3.2.3 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#

```

### 5.3.3 Default Network Queue Policy Values

The default network queue policy is named **default** and cannot be modified or deleted. The default policy parameters are shown below.

```

*A:PE>config>qos# network-queue "default"
*A:PE>config>qos>network-queue# info detail

description "Default network queue QoS policy."
hs-attachment-policy "default"
queue 1 auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 0 fir 0
 adaptation-rule pir closest cir closest fir closest

```

```
mbs 50
cbs 1
hs-mbs 100
hs-wrr-weight 1
hs-class-weight 1
hs-wred-queue policy "_tmnx_hs_default"
no hs-alt-port-class-pool
no pool
drop-tail
 low
 percent-reduction-from-mbs default
 exit
exit
exit
queue 2 auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 25 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 50
 cbs 3
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
 exit
exit
queue 3 auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 25 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 50
 cbs 10
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
 exit
exit
queue 4 auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 25 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 25
 cbs 3
```

```
hs-mbs 100
hs-wrr-weight 1
hs-class-weight 1
hs-wred-queue policy "_tmnx_hs_default"
no hs-alt-port-class-pool
no pool
drop-tail
 low
 percent-reduction-from-mbs default
 exit
exit
exit
queue 5 auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 100 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 50
 cbs 10
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
 exit
exit
queue 6 auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 100 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 50
 cbs 10
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
 exit
exit
queue 7 auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 10 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 25
 cbs 3
 hs-mbs 100
 hs-wrr-weight 1
```

```
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
 exit
exit
queue 8 auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 10 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 25
 cbs 3
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
exit
exit
queue 9 multipoint auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 0 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 50
 cbs 1
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
exit
exit
queue 10 multipoint auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 5 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 50
 cbs 1
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
```



```

 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
 exit
exit
queue 11 multipoint auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 5 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 50
 cbs 1
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
exit
queue 12 multipoint auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 5 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 25
 cbs 1
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
exit
queue 13 multipoint auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 100 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 50
 cbs 1
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool

```

```
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
exit
queue 14 multipoint auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 100 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 50
 cbs 1
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
exit
queue 15 multipoint auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 10 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 25
 cbs 1
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
 percent-reduction-from-mbs default
 exit
exit
queue 16 multipoint auto-expedite create
 no port-parent
 no avg-frame-overhead
 rate 100 cir 10 fir 0
 adaptation-rule pir closest cir closest fir closest
 mbs 25
 cbs 1
 hs-mbs 100
 hs-wrr-weight 1
 hs-class-weight 1
 hs-wred-queue policy "_tmnx_hs_default"
 no hs-alt-port-class-pool
 no pool
 drop-tail
 low
```

```
 percent-reduction-from-mbs default
 exit
exit
egress-hsmda
 no wrr-policy
 no packet-byte-offset
 queue 1
 adaptation-rule pir closest
 percent-rate 100.00
 wrr-weight 1
 mbs default
 no burst-limit
 slope-policy "default"
 exit
 queue 2
 adaptation-rule pir closest
 percent-rate 100.00
 wrr-weight 1
 mbs default
 no burst-limit
 slope-policy "default"
 exit
 queue 3
 adaptation-rule pir closest
 percent-rate 100.00
 wrr-weight 1
 mbs default
 no burst-limit
 slope-policy "default"
 exit
 queue 4
 adaptation-rule pir closest
 percent-rate 100.00
 no wrr-weight
 mbs default
 no burst-limit
 slope-policy "default"
 exit
 queue 5
 adaptation-rule pir closest
 percent-rate 100.00
 no wrr-weight
 mbs default
 no burst-limit
 slope-policy "default"
 exit
 queue 6
 adaptation-rule pir closest
 percent-rate 100.00
 no wrr-weight
 mbs default
 no burst-limit
 slope-policy "default"
 exit
 queue 7
 adaptation-rule pir closest
 percent-rate 100.00
 no wrr-weight
```

```
 mbs default
 no burst-limit
 slope-policy "default"
 exit
 queue 8
 adaptation-rule pir closest
 percent-rate 100.00
 no wrr-weight
 mbs default
 no burst-limit
 slope-policy "default"
 exit
exit
hs-wrr-group 1
 rate 100
 adaptation-rule pir closest
 hs-class-weight 1
exit
hs-wrr-group 2
 rate 100
 adaptation-rule pir closest
 hs-class-weight 1
exit
fc af create
 multicast-queue 11
 queue 3
 egress-hsmda
 queue 3
 exit
exit
fc be create
 multicast-queue 9
 queue 1
 egress-hsmda
 queue 1
 exit
exit
fc ef create
 multicast-queue 14
 queue 6
 egress-hsmda
 queue 6
 exit
exit
fc h1 create
 multicast-queue 15
 queue 7
 egress-hsmda
 queue 7
 exit
exit
fc h2 create
 multicast-queue 13
 queue 5
 egress-hsmda
 queue 5
 exit
exit
fc l1 create
```

```
 multicast-queue 12
 queue 4
 egress-hsmda
 queue 4
 exit
 exit
fc l2 create
 multicast-queue 10
 queue 2
 egress-hsmda
 queue 2
 exit
exit
fc nc create
 multicast-queue 16
 queue 8
 egress-hsmda
 queue 8
 exit
exit

*A:PE>config>qos>network-queue#
```

---

## 5.4 Service Management Tasks

This section discusses network queue QoS policy service management tasks.

### 5.4.1 Deleting QoS Policies

A network queue policy is associated by default with XMA or MDA network ingress ports. The default policy can be replaced with a customer-configured policy, but cannot entirely be removed. When a QoS policy is removed, the policy association reverts to the default network-queue policy.

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`

### 5.4.2 Removing 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`

### 5.4.3 Copying and Overwriting QoS Policies

An existing network queue policy can be copied, renamed with a new policy ID name, or used to 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`

## **5.4.4 Editing QoS Policies**

Existing policies, except the default policies, and entries in the CLI can be modified. 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, then overwrite the original policy.





## 5.5 Network Queue QoS Policy Command Reference

### 5.5.1 Command Hierarchies

- [Configuration Commands](#)
- [Operational Commands](#)
- [Show Commands](#)

#### 5.5.1.1 Configuration Commands

```

config
 — qos
 — network-queue policy-name [create]
 — no network-queue policy-name
 — description description-string
 — no description
 — egress-hsmda
 — packet-byte-offset {add add-bytes | subtract sub-bytes}
 — no packet-byte-offset
 — [no] queue queue-id
 — adaptation-rule [pir {min | max | closest}]
 — no adaptation-rule
 — burst-limit size [bytes | kilobytes]
 — no burst-limit
 — mbs size [bytes | kilobytes]
 — no mbs
 — percent-rate percentage
 — no percent-rate
 — slope-policy hsmda-slope-policy-name
 — no slope-policy
 — wrr-weight value
 — no wrr-weight
 — wrr-policy hsmda-wrr-policy-name
 — no wrr-policy
 — fc fc-name [create]
 — no fc fc-name
 — egress-hsmda
 — queue queue-id
 — no queue
 — multicast-queue queue-id
 — no multicast-queue
 — queue queue-id
 — no queue

```

- 
- **queue** *queue-id* [*multipoint*] [*queue-type*] [*create*]
  - **no queue** *queue-id*
    - **adaptation-rule** [*pir* {*max* | *min* | *closest*}] [*cir* {*max* | *min* | *closest*}] [*fir* {*max* | *min* | *closest*}]
    - **no adaptation-rule**
    - **avg-frame-overhead** *percent*
    - **no avg-frame-overhead**
    - **cbs** *percent*
    - **no cbs**
    - **drop-tail**
      - **low**
        - **percent-reduction-from-mbs** *percent*
        - **no percent-reduction-from-mbs**
    - **mbs** *percent*
    - **no mbs**
    - **pool** *pool-name*
    - **no pool**
    - **port-parent** [*weight weight*] [*level level*] [*cir-weight cir-weight*] [*cir-level cir-level*]
    - **no port-parent**
    - **rate** *percent* [*cir percent*] [*fir percent*]
    - **no rate**

### 5.5.1.2 Operational Commands

- ```

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

```

5.5.1.3 Show Commands

- ```

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

```

## 5.5.2 Command Descriptions

### 5.5.2.1 Configuration Commands

#### 5.5.2.1.1 Generic Commands

#### description

|                    |                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>description</b> <i>description-string</i><br><b>no description</b>                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>qos>network-queue                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The <b>description</b> command associates a text string with a configuration context to help identify the context in the configuration file.</p> <p>The <b>no</b> form of this command removes any description string from the context.</p> |
| <b>Default</b>     | no description                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.                                                           |

#### 5.5.2.1.2 Network Queue QoS Policy Commands

#### network-queue

|                    |                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>network-queue</b> <i>policy-name</i> [create]<br><b>no network-queue</b> <i>policy-name</i>                                                                                                                                                         |
| <b>Context</b>     | config>qos                                                                                                                                                                                                                                             |
| <b>Description</b> | This command creates a context to configure a network queue policy. Network queue policies define the ingress network queuing at the XMA or MDA network node level and on the Ethernet port and SONET/SDH path level to define network egress queuing. |

---

|                   |                                                                                                                                                                                                                                                                                                                                    |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | network-queue "default"                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b> | <i>policy-name</i> — The name of the network queue policy.                                                                                                                                                                                                                                                                         |
|                   | <p><b>Values</b> Valid names consist of any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><b>create</b> — Required keyword when creating a network queue policy.</p> |

## egress-hsmda

|                    |                                                                                   |
|--------------------|-----------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>egress-hsmda</b>                                                               |
| <b>Context</b>     | config>qos>network-queue<br>config>qos>network-queue>fc                           |
| <b>Description</b> | This command enters the context to configure queue definitions for use on HSMDAs. |

## packet-byte-offset

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>packet-byte-offset</b> { <b>add</b> <i>add-bytes</i>   <b>subtract</b> <i>sub-bytes</i> }<br><b>no packet-byte-offset</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>qos>network-queue>egress-hsmda                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>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 14-byte Ethernet DLC header, 4-byte or 8-byte VLAN tag (optional), 20-byte IP header, IP payload, and the 4-byte CRC (everything except the preamble and inter-frame gap). For example, the <b>packet-byte-offset</b> command can be used to add the frame encapsulation overhead (20 bytes) to the queue's accounting functions. The accounting functions affected include:</p> <ul style="list-style-type: none"> <li>• Offered High Priority / In-Profile Octet Counter</li> <li>• Offered Low Priority / Out-of-Profile Octet Counter</li> <li>• Discarded High Priority / In-Profile Octet Counter</li> <li>• Discarded Low Priority / Out-of-Profile Octet Counter</li> <li>• Forwarded In-Profile Octet Counter</li> <li>• Forwarded Out-of-Profile Octet Counter</li> <li>• Peak Information Rate (PIR) Leaky Bucket Updates</li> <li>• Committed Information Rate (CIR) Leaky Bucket Updates</li> <li>• Queue Group Aggregate Rate Limit Leaky Bucket Updates</li> </ul> |

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 **add** or **subtract** keywords 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. For example, in 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.

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 subscriber packets on an Ethernet aggregation network.

The **packet-byte-offset** value may be overridden for the HSMDA queue at the network queue level.

The **no** form of this command removes any accounting size changes to packets handled by the queue.

|                   |                                                                                                                                                                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <b>add</b> <i>add-bytes</i> — Indicates that the byte value should be added to the packet for queue and queue group level accounting functions.                                                                                             |
|                   | <b>Values</b> 1 to 31                                                                                                                                                                                                                       |
|                   | <b>subtract</b> <i>sub-bytes</i> — Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The minimum resulting packet size used by the system is 64 bytes with an HSMDA. |
|                   | <b>Values</b> 1 to 64                                                                                                                                                                                                                       |

## queue

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] queue</b> <i>queue-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>qos>network-queue>egress-hsmda                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command 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 create or dynamically assign the queue to the object to which the policy is applied. The queue identified by <i>queue-id</i> always exists whether the command is executed or not. For HSMDA SAPs and subscribers, all eight queues exist when the system allocates an HSMDA queue group to the object. |

With standard service queues, the scheduling behavior relative to other queues is based on two items: the queue's Best-Effort or Expedited nature and the dynamic rate of the queue relative to the defined CIR. The HSMMDA queues are handled differently. The create time auto-expedite, explicit expedite, and best-effort qualifiers have been eliminated; instead, the scheduling behavior is based solely on the queues identifier. Queues with a *queue-id* of 1 are placed in scheduling class 1. Queues with a *queue-id* of 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.

Unlike standard service queues, the HSMMDA 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 dependent on the ingress sub-forwarding class to which the packet is mapped.

The **no** form of this command restores the defined *queue-id* to its default parameters. All HSMMDA queues having the queue-id and associated with the QoS policy are re-initialized to default parameters.

**Parameters**     *queue-id* — Defines the context from which the eight egress queues will be entered for editing purposes.

## adaptation-rule

**Syntax**     **adaptation-rule** [**pir** {**max** | **min** | **closest**}]  
**no adaptation-rule**

**Context**     config>qos>network-queue>egress-hsmda>queue

**Description**     This command defines the method used by the system to derive the operational PIR settings when the HSMMDA queue is provisioned in hardware. For the PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of this command removes any explicitly defined constraints used to derive the operational PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **pir** apply.

**Parameters**     **pir** — Defines the constraints enforced when adapting the PIR rate defined. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the HSMMDA queue. When the **pir** command is not specified, the default applies.

- max** — Specifies that the operational PIR for the HSMDA queue will be less than or equal to the rate specified using the **rate** command.
- min** — Specifies that the operational PIR for the HSMDA queue will be greater than or equal to the rate specified using the **rate** command.
- closest** — Specifies that the operational PIR for the HSMDA queue will be the rate closest to the rate specified using the **rate** command.
- Default**      closest

## burst-limit

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>burst-limit</b> size [bytes   kilobytes]<br><b>no burst-limit</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>qos>network-queue>egress-hsmda>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command defines 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.</p> <p>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.</p> <p>The <b>no</b> form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying the 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.</p>                                                                                             |
| <b>Parameters</b>  | <p><b>size</b> — 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.</p> <p><b>Values</b>      1 to 14,000 (14,000 or 14,000,000 depending on bytes or kilobytes)</p> <p><b>bytes</b> — Specifies 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.</p> <p><b>kilobytes</b> — Specifies 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.</p> |

## mbs

|               |                                                     |
|---------------|-----------------------------------------------------|
| <b>Syntax</b> | <b>mbs</b> size [byte   kilobytes]<br><b>no mbs</b> |
|---------------|-----------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>qos>network-queue>egress-hsmda>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>The Maximum Burst Size (<b>mbs</b>) command specifies the relative amount of the buffer pool space for the maximum buffers for a specific HSMDA queue.</p> <p>The MBS value is used by a queue to determine whether it has exhausted its total allowed buffers while enqueueing packets. When 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 against 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.</p> <p>The <b>no</b> form of the command returns the MBS size for the queue to the default for the forwarding class.</p> |
| <b>Parameters</b>  | <p><b>size</b> — Specifies the size of the MBS for the HSMDA queue.</p> <p><b>Values</b>      0 to 2688000 (for bytes)<br/>                  0 to 2625 (for kilobytes)</p> <p><b>bytes</b> — Identifies the size value in terms of bytes.</p> <p><b>kilobytes</b> — Identifies the size value in terms of kilobytes.</p> <p><b>default</b> — Specifies the size of the MBS to be the default value configured on the system.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## percent-rate

|                    |                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>percent-rate</b> <i>percentage</i><br><b>no percent-rate</b>                                                                                                            |
| <b>Context</b>     | config>qos>network-queue>egress-hsmda>queue                                                                                                                                |
| <b>Description</b> | <p>This command specifies the PIR shaping rate for the HSMDA queue.</p> <p>The <b>no</b> form of this command returns the PIR size for the queue to the default value.</p> |
| <b>Parameters</b>  | <p><b>percentage</b> — Specifies the PIR percentage rate for the HSMDA queue.</p> <p><b>Values</b>      0.10 to 100.00</p>                                                 |

## slope-policy

|                |                                                                              |
|----------------|------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>slope-policy</b> <i>hsmda-slope-policy-name</i><br><b>no slope-policy</b> |
| <b>Context</b> | config>qos>network-queue>egress-hsmda>queue                                  |



**Description** This command associates an existing HSMDB slope policy with the QoS policy HSMDB queue. The specified hsmdb-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. When a slope policy is associated with a QoS policy queue or 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 a 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 HSMDB slope policies may be configured on a system.

The system maintains a slope policy named hsmdb-default, which acts as a default policy when an explicit slope policy has not been defined for an HSMDB queue. The default policy may be edited, but it cannot be deleted. If a no slope-policy hsmdb-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 this 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.

|                   |                                                                                                                                                                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>hsmda-slope-policy-name</i> — 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. When a slope policy is associated with an HSMDA queue, the policy cannot be deleted. |
| <b>Default</b>    | hsmda-default                                                                                                                                                                                                                                                                                                               |

wrr-weight

|                    |                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>wrr-weight</b> <i>value</i><br><b>no wrr-weight</b>                                                                                                           |
| <b>Context</b>     | config>qos>network-queue>egress-hsmda>queue                                                                                                                      |
| <b>Description</b> | This command assigns the weight value to the HSMDA queue.<br><br>The <b>no</b> form of this command returns the weight value for the queue to the default value. |
| <b>Parameters</b>  | <i>value</i> — Specifies the weight for the HSMDA queue.                                                                                                         |
| <b>Values</b>      | 1 to 32                                                                                                                                                          |

wrr-policy

|               |                                                                        |
|---------------|------------------------------------------------------------------------|
| <b>Syntax</b> | <b>wrr-policy</b> <i>hsmda-wrr-policy-name</i><br><b>no wrr-policy</b> |
|---------------|------------------------------------------------------------------------|

---

|                    |                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>qos>network-queue>egress-hsmda                                                                             |
| <b>Description</b> | This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy with the HSMDA queue. |
| <b>Parameters</b>  | <i>hsmda-wrr-policy-name</i> — Specifies the existing HSMDA WRR policy name to associate with the queue.          |

## fc

|                    |                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fc</b> <i>fc-name</i> [ <b>create</b> ]<br><b>no fc</b> <i>fc-name</i>                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>network-queue                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | The <b>fc</b> command is used to enter the forwarding class mapping context for the given <i>fc-name</i> . Each forwarding class maps by default to queues 1 (unicast) and 9 (multipoint).                                                                                                                             |
| <b>Parameters</b>  | <i>fc-name</i> — A valid forwarding class must be specified as <i>fc-name</i> when the <b>fc</b> command is executed. When the <b>fc</b> <i>fc-name</i> command is successfully executed, the system will enter the specified forwarding class context where the <b>queue</b> <i>queue-id</i> command may be executed. |
| <b>Values</b>      | be, l2, af, l1, h2, ef, h1, nc                                                                                                                                                                                                                                                                                         |
| <b>create</b>      | — Required parameter when creating an FC node.                                                                                                                                                                                                                                                                         |

## queue

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>queue</b> <i>queue-id</i><br><b>no queue</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>qos>network-queue>fc<br>config>qos>network-queue>fc>egress-hsmda                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | This command 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 create or dynamically assign the queue to the object to which the policy is applied. The queue identified by <i>queue-id</i> always exists whether the command is executed or not. In the case of HSMDA SAPs and subscribers, all eight queues exist at the moment the system allocates an HSMDA queue group to the object. |

### 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; 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.

## Every HSMDB Queue Supports Profile Mode Implicitly

Unlike standard service queues, the HSMMDA 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 dependent on the ingress sub-forwarding class to which the packet is mapped.

The **no** form of this command restores the defined *queue-id* to its default parameters. All HSMDBA queues having the *queue-id* and associated with the QoS policy are re-initialized to default parameters.

**Parameters** *queue-id* — Defines the context from which the eight egress queues will be entered for editing purposes.

## multicast-queue

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

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

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

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 this 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 XMA or 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.

|                   |                                                                                                                                                                                                                               |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>queue-id</i> — Specifies any valid multipoint queue-ID in the policy. The <i>queue-id</i> parameter specified must be an existing, multipoint queue defined in the <b>config&gt;qos&gt;network-queue&gt;queue</b> context. |
| <b>Values</b>     | 1 to 16                                                                                                                                                                                                                       |
| <b>Default</b>    | 11                                                                                                                                                                                                                            |

**Network Queue QoS Policy Queue Commands**

queue

|                    |                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>queue</b> <i>queue-id</i> [ <b>multipoint</b> ] [ <i>queue-type</i> ] [ <b>create</b> ]<br><b>no queue</b> <i>queue-id</i> |
| <b>Context</b>     | config>qos>network-queue                                                                                                      |
| <b>Description</b> | This command enters the context to configure a QoS network-queue policy queue.                                                |

Explicit definition of an ingress queue's type status is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or best effort 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 best effort 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 queue type 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 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 a QoS policy with multipoint queues is applied to an Epipe or IES SAP, the multipoint queues are not created. 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 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 to 32

**multipoint** — This optional keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be used to forward multicast, broadcast, or unknown unicast ingress traffic.

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 preexisting multipoint queue to edit *queue-id* parameters.

**Default**            Non-multipoint (unicast queue)

*queue-type* — The **expedite**, **best-effort**, and **auto-expedite** queue types are mutually exclusive. 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.

**Values**            expedite, best-effort, auto-expedite

**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

## adaptation-rule

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>adaptation-rule</b> [ <b>pir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }] [ <b>cir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }] [ <b>fir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }]<br><b>no adaptation-rule</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>qos>network-queue>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command defines the method used by the system to derive the operational FIR, CIR, and PIR settings when the queue is provisioned in hardware. For the FIR, CIR, and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.</p> <p>When configured on an egress HSQ queue group queue, the <b>cir</b> keyword is ignored. This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the <b>adaptation-rule</b> is performed under the <b>hs-wrr-group</b> within the network queue policy.</p> <p>The <b>no</b> form of this command removes any explicitly defined constraints used to derive the operational FIR, CIR, and PIR created by the application of the policy. When a specific <b>adaptation-rule</b> is removed, the default constraints for <b>fir</b>, <b>cir</b>, and <b>pir</b> apply.</p> |
| <b>Default</b>     | adaptation-rule pir closest cir closest fir closest                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <p><b>pir</b> — Defines the constraints enforced when adapting the queue's PIR. The <b>pir</b> parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the <b>pir</b> parameter is not specified, the default applies.</p> <p><b>cir</b> — Defines the constraints enforced when adapting the queue's CIR. The <b>cir</b> parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the <b>cir</b> parameter is not specified, the default constraint applies.</p>                                                                                                                                                                                                                                                                                                                                                                                                            |

**fir** — Defines the constraints enforced when adapting the queue's FIR. The **fir** parameter requires a qualifier that defines the constraint used when deriving the operational FIR for the queue. When the **fir** parameter is not specified, the default constraint applies. FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

**max** — Specifies that the operational rate for the queue will be equal to or less than the requested rate.

**min** — Specifies that the operational rate for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

**closest** — Specifies that the operational rate for the queue will be the rate closest to the requested rate.

## avg-frame-overhead

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>avg-frame-overhead</b> <i>percent</i><br><b>no avg-frame-overhead</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>qos>network-queue>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | 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 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 1000 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, the frame-encapsulation overhead would be 50 x 20 or 1000 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 1000 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 1000 octets and the offered-load is 10 000 octets, then the packet to frame factor would be  $1000 / 10\,000$  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, 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  $500 \times 1.1$  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 11 000 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, 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  $7500 \times 1.1$  or 8250 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 (applies only to the 7450 ESS and 7750 SR) — 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%. 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.

|                   |                                                                                                                                                                                            |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no avg-frame-overhead                                                                                                                                                                      |
| <b>Parameters</b> | <i>percent</i> — 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. |
| <b>Values</b>     | 0.00 to 100.00                                                                                                                                                                             |

## cbs

|                    |                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cbs</b> <i>percent</i><br><b>no cbs</b>                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>network-queue>queue                                                                                                                                                                                                                          |
| <b>Description</b> | The Committed Burst Size ( <b>cbs</b> ) command specifies the relative amount of reserved buffers for a specific ingress network XMA or 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 enqueueing packets. When 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. At egress, there are two additional RED slopes maintained in each buffer pool: the highplus slope is used by inplus-profile packets, and the exceed slope is used by exceed-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 XMA or 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.

**Default** The cbs forwarding class defaults are listed in the [Table 25](#).

**Table 25 CBS Forwarding Class Defaults**

| Forwarding Class | Forwarding 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           |

**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 when 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. Forwarding class queue on egress channel applies only to the 7450 ESS and 7750 SR.

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 XMA or MDA** — For network ingress, each forwarding class is supported by an ingress queue per XMA or MDA. These forwarding class queues are automatically created when a single port or channel is placed in the network mode on the XMA/MDA and are removed when all network ports or channels are removed from the XMA/MDA (defined as access). The configuration parameters for each queue come from the applied ingress policy under the network context of the XMA or 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 XMA or MDA. This buffer pool will dynamically fluctuate based on the sum of all ingress pool sizes for all network ports and channels on the XMA or 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.

**Parameters**     *percent* — The percent of buffers reserved from the total buffer pool space, expressed as a decimal integer. If 10 Mbytes is the total buffer space in the buffer pool, a value of 10 would reserve 1 Mbyte (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 to 100

## drop-tail

**Syntax**        **drop-tail**

**Context**        config>qos>network-queue>queue

**Description**    This command enters the context to configure queue drop tail parameters.

## low

**Syntax**        **low**

**Context**        config>qos>network-queue>queue>drop-tail

**Description**    This command enters the context to configure the queue low drop tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets will not be accepted into the queue and will be discarded.

## percent-reduction-from-mbs

**Syntax**        **percent-reduction-from-mbs** *percent*

### no percent-reduction-from-mbs

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>qos>network-queue>queue>drop-tail>low                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command configures the ingress and egress network queue low drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and <b>percent-reduction-from-mbs</b> is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value and will be discarded.</p> <p>The exceed drop tail is not configurable for network queues, however, it is set to a value of 10% in addition to low drop tail and capped by the MBS.</p> |
| <b>Default</b>     | percent-reduction-from-mbs 10                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>percent</i> — Specifies the percentage reduction from the MBS for a queue drop tail.</p> <p><b>Values</b> 0 to 100, default</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## mbs

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>        | <p><b>mbs percent</b></p> <p><b>no mbs</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>       | config>qos>network-queue>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>   | <p>This command specifies the relative amount of buffer pool space for the maximum buffers for a specific ingress network XMA or MDA forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.</p> <p>The MBS value is used by a queue to determine whether it has exhausted its total allowed buffers while enqueueing packets. When 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 is not guaranteed to have a buffer available when needed or that the packet's RED slope will not force the discard of the packet. In order to safeguard against queue starvation (when a queue does not receive its fair share of buffers), set proper CBS parameters and control CBS oversubscription. Another safeguard is to properly set the RED slope parameters for the needs of the network queues.</p> <p>The MBS 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.</p> <p>The <b>no</b> form of this command returns the MBS for the queue to the default for the forwarding class.</p> |
| <b>Special Cases</b> | <p><b>Forwarding Class Queue on Egress Network Port or Channel</b> — 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 when a port or channel is placed in network mode. The configuration parameters for each queue come from the applied egress policy on the network port or channel. Forwarding</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

class queue on egress channel applies only to the 7450 ESS and 7750 SR.

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 over-subscription 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 XMA or MDA** — For network ingress, each forwarding class is supported by an ingress queue per XMA or MDA. These forwarding class queues are automatically created when a single port or channel is placed in the network mode on the XMA/MDA and are removed when all network ports or channels are removed from the XMA/MDA (defined as access). The configuration parameters for each queue come from the applied ingress policy under the network context of the XMA or 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 XMA or MDA. This buffer pool will dynamically fluctuate based on the sum of all ingress pool sizes for all network ports and channels on the XMA or 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 over-subscription 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.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>percent</i> — The percent of buffers from the total buffer pool space for the maximum amount of buffers, expressed as a decimal integer. If 10 Mbytes is the total buffer space in the buffer pool, a value of 10 would limit the maximum queue size to 1 Mbyte (10%) of buffer space for the forwarding class queue. If the total size is increased to 20 Mbytes, the existing value of 10 would automatically increase the maximum size of the queue to 2 Mbytes. |
| <b>Values</b>     | 0 to 100                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

pool

|                |                                                |
|----------------|------------------------------------------------|
| <b>Syntax</b>  | <b>pool</b> <i>pool-name</i><br><b>no pool</b> |
| <b>Context</b> | config>qos>network-queue>queue                 |

**Description** This command configures a pool for the network queue within the specified 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) in the **show queue** command output using the **pool** keyword.

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. Named pools are not supported on the 7950 XRS, 7750 SR-c4/12, 7750 SR-a4/a8, or 7750 SR-1e/2e/3e.

The **no** form of this command removes a named pool association for the queue. When the pool name is removed, the queue will be placed on the appropriate default pool.

**Default** no pool

**Parameters** *pool-name* — Identifies a named pool, up to 16 characters, where the policy will be applied.

## 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. When 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 XMA, MDA, or CMA for ingress queue definition, the port-parent association of the queues is 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 then the queue or scheduler instance exists, the queue or scheduler will become orphaned.

**Default** no port-parent

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <b>weight</b> <i>weight</i> — Defines the weight the queue or scheduler will use at the above-CIR port priority level (defined by the level parameter).                                                                                                                                                                                                                                                                                                                                                                                     |
|                   | <b>Values</b> 0 to 100                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|                   | <b>Default</b> 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                   | <b>level</b> <i>level</i> — Defines the port priority the queue or scheduler will use to receive bandwidth for its above-CIR offered-load.                                                                                                                                                                                                                                                                                                                                                                                                  |
|                   | <b>Values</b> 1 to 8 (8 is the highest priority)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                   | <b>Default</b> 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                   | <b>cir-weight</b> <i>cir-weight</i> — 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 is used. |
|                   | <b>Values</b> 0 to 100                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|                   | <b>Default</b> 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                   | <b>cir-level</b> <i>cir-level</i> — 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.                                                                                     |
|                   | <b>Values</b> 0 to 8 (8 is the highest priority)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                   | <b>Default</b> 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## rate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rate</b> <i>percent</i> [ <b>cir</b> <i>percent</i> ] [ <b>fir</b> <i>percent</i> ]<br><b>no rate</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>qos>network-queue>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command defines the administrative Peak Information Rate (PIR), the administrative Committed Information Rate (CIR), and the administrative Fair Information Rate (FIR) parameters for the queue.</p> <p>The PIR defines the percentage that the queue can transmit packets through the switch fabric (for ingress queues) or out of an egress port (for 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 bandwidth.</p> <p>The CIR defines the percentage at which the system prioritizes the queue over other queues competing for the same bandwidth.</p> |



The CIR can be used by the queue's **port-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 port scheduler.

The FIR defines an additional percentage at which the system prioritizes the queue over other queues competing for the same bandwidth above that used by the CIR percentage.

The **rate** command can be executed at any time, altering the PIR, CIR, and FIR for all queues created through the association of the network queue policy with the *queue-id*.

When configured on an egress HSQ queue group queue, the **cir** keyword is ignored.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the rate is performed under the **hs-wrr-group** within the network queue policy.

The **no** form of the command returns all queues created with the *queue-id* by association with the network queue policy to the default PIR, CIR, and FIR parameters.

**Default**      rate 100 cir 0 fir 0

**Parameters**    *percent* — Defines the percentage of the FP ingress capacity or egress port speed for the guaranteed rate allowed for the queue. When the **rate** command is executed, a valid *percent* (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 the value must be given as a positive integer.

The actual PIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values**        1 to 100

**Default**        100

**cir percent** — Defines the percentage of the FP ingress capacity or egress port speed for the CIR 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 the value must be given as a positive integer. The actual CIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values**        0 to 100

**Default**        0

**fir percent** — Defines the percentage of the FP ingress capacity or egress port speed for the FIR allowed for the queue. When the **rate** command is executed, a FIR setting is optional. When the **rate** command has not been executed or the **fir** parameter is not explicitly specified, the default FIR (0) is assumed. Fractional values are not allowed and the value must be given as a positive integer. The actual FIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned. FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

**Values** 0 to 100

**Default** 0

### 5.5.2.2 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.

```
SR7>config>qos# copy network-queue nq1 nq2
MINOR: CLI Destination "nq2" exists - use {overwrite}.
SR7>config>qos# copy network-queue nq1 nq2 overwrite
```

### 5.5.2.3 Show Commands

#### network-queue

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>network-queue</b> [ <i>network-queue-policy-name</i> ] [ <b>detail</b> ]                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | show>qos                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command displays network queue policy information.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <p><i>network-queue-policy-name</i> — The name of the network queue policy.</p> <p><b>Values</b> Valid names consist of any string, up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.</p> <p><b>detail</b> — Includes each queue's rates and adaptation-rule and cbs details. It also shows FC-to-queue mapping details.</p> |
| <b>Output</b>      | The following output is an example of network-queue information.                                                                                                                                                                                                                                                                                                                                                                                                              |

#### Sample Output

```
*A:PE# show qos network-queue
=====
Network Queue Policies
=====
Policy-Id Description

default Default network queue QoS policy.
=====
*A:PE#
*A:PE# show qos network-queue "default"
=====
QoS Network Queue Policy
=====

Network Queue Policy (default)

Policy : default
Description : Default network queue QoS policy.
Wrr Policy : (Not Specified)
Pkt.Byte Offset : 0
HS Attachment Policy: default

HSMDA Queues

Queue PIR PIR-Rule Wrr-Wght MBS Slope Plcy Burst Lmt

1 def closest def def default default
2 def closest def def default default
3 def closest def def default default
```

```

4 def closest def def default default
5 def closest def def default default
6 def closest def def default default
7 def closest def def default default
8 def closest def def default default

```

-----  
Associations  
-----

```

MDA : 1/1
MDA : 1/2
Port-id : 1/1/2
Port-id : 1/1/3
Port-id : 1/1/4
Port-id : 1/1/5
Port-id : 1/1/6
Port-id : 1/1/7
Port-id : 1/1/8

```

```

=====
*A:PE#
*A:PE# show qos network-queue "default" detail
=====

```

QoS Network Queue Policy  
-----

Network Queue Policy (default)  
-----

```

Policy : default
Description : Default network queue QoS policy.
Wrr Policy : (Not Specified)
Pkt.Byte Offset : 0
HS Attachment Policy: default

```

Queue Information  
-----

|                    |                   |                  |                 |
|--------------------|-------------------|------------------|-----------------|
| Queue-Id           | : 1               | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 0             |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 50.00           | CBS              | : 1.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 2               | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 25            |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 50.00           | CBS              | : 3.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 3               | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 25            |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |

|                    |                   |                  |                 |
|--------------------|-------------------|------------------|-----------------|
| MBS                | : 50.00           | CBS              | : 10.00         |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 4               | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 25            |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 25.00           | CBS              | : 3.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 5               | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 100           |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 50.00           | CBS              | : 10.00         |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 6               | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 100           |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 50.00           | CBS              | : 10.00         |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 7               | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 10            |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 25.00           | CBS              | : 3.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 8               | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 10            |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 25.00           | CBS              | : 3.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 9               | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 0             |
| PIR Rule           | : closest         | CIR Rule         | : closest       |

---

|                    |                   |                  |                 |
|--------------------|-------------------|------------------|-----------------|
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 50.00           | CBS              | : 1.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 10              | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 5             |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 50.00           | CBS              | : 1.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 11              | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 5             |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 50.00           | CBS              | : 1.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 12              | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 5             |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 25.00           | CBS              | : 1.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 13              | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 100           |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 50.00           | CBS              | : 1.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 14              | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 100           |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 50.00           | CBS              | : 1.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 15              | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 10            |

|                    |                   |                  |                 |
|--------------------|-------------------|------------------|-----------------|
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 25.00           | CBS              | : 1.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |
| Queue-Id           | : 16              | Queue-Type       | : auto-expedite |
| PIR                | : 100             | CIR              | : 10            |
| PIR Rule           | : closest         | CIR Rule         | : closest       |
| FIR                | : 0               | FIR Rule         | : closest       |
| MBS                | : 25.00           | CBS              | : 1.00          |
| Avg Frame Overhead | : 0.00            |                  |                 |
| Named-Buffer Pool  | : (Not Specified) |                  |                 |
| Parent             | : No              |                  |                 |
| PIR Level/Weight   | : n/a             | CIR Level/Weight | : n/a           |
| Low Drop Tail      | : def             |                  |                 |

-----  
HS Queue Information  
-----

|                      |                    |               |          |
|----------------------|--------------------|---------------|----------|
| Queue-Id             | : 1                |               |          |
| PIR                  | : 100              |               |          |
| PIR Rule             | : closest          |               |          |
| HS Class Weight      | : 1                | HS Wrr Weight | : 1      |
| HS Wred-Q Slope Plcy | : _tmnx_hs_default |               |          |
| HS Alt Class Pool    | : False            | HS MBS        | : 100.00 |
| Queue-Id             | : 2                |               |          |
| PIR                  | : 100              |               |          |
| PIR Rule             | : closest          |               |          |
| HS Class Weight      | : 1                | HS Wrr Weight | : 1      |
| HS Wred-Q Slope Plcy | : _tmnx_hs_default |               |          |
| HS Alt Class Pool    | : False            | HS MBS        | : 100.00 |
| Queue-Id             | : 3                |               |          |
| PIR                  | : 100              |               |          |
| PIR Rule             | : closest          |               |          |
| HS Class Weight      | : 1                | HS Wrr Weight | : 1      |
| HS Wred-Q Slope Plcy | : _tmnx_hs_default |               |          |
| HS Alt Class Pool    | : False            | HS MBS        | : 100.00 |
| Queue-Id             | : 4                |               |          |
| PIR                  | : 100              |               |          |
| PIR Rule             | : closest          |               |          |
| HS Class Weight      | : 1                | HS Wrr Weight | : 1      |
| HS Wred-Q Slope Plcy | : _tmnx_hs_default |               |          |
| HS Alt Class Pool    | : False            | HS MBS        | : 100.00 |
| Queue-Id             | : 5                |               |          |
| PIR                  | : 100              |               |          |
| PIR Rule             | : closest          |               |          |
| HS Class Weight      | : 1                | HS Wrr Weight | : 1      |
| HS Wred-Q Slope Plcy | : _tmnx_hs_default |               |          |
| HS Alt Class Pool    | : False            | HS MBS        | : 100.00 |
| Queue-Id             | : 6                |               |          |
| PIR                  | : 100              |               |          |
| PIR Rule             | : closest          |               |          |
| HS Class Weight      | : 1                | HS Wrr Weight | : 1      |
| HS Wred-Q Slope Plcy | : _tmnx_hs_default |               |          |
| HS Alt Class Pool    | : False            | HS MBS        | : 100.00 |
| Queue-Id             | : 7                |               |          |

---

```

PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00
Queue-Id : 8
PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00
Queue-Id : 9
PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00
Queue-Id : 10
PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00
Queue-Id : 11
PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00
Queue-Id : 12
PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00
Queue-Id : 13
PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00
Queue-Id : 14
PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00
Queue-Id : 15
PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00
Queue-Id : 16
PIR : 100
PIR Rule : closest
HS Class Weight : 1 HS Wrr Weight : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False HS MBS : 100.00

```



```

HS Wrr Group Information

HS Wrr Group Id : 1 Class Weight : 1
Percent Rate : 100 PIR Adaptation Rule: closest
HS Wrr Group Id : 2 Class Weight : 1
Percent Rate : 100 PIR Adaptation Rule: closest

HSMDBA Queues

Queue PIR PIR-Rule Wrr-Wght MBS Slope Plcy Burst Lmt

1 def closest def def default default
2 def closest def def default default
3 def closest def def default default
4 def closest def def default default
5 def closest def def default default
6 def closest def def default default
7 def closest def def default default
8 def closest def def default default

FC UCastQ MCastQ EHsmdaQ

be 1 9 1
l2 2 10 2
af 3 11 3
l1 4 12 4
h2 5 13 5
ef 6 14 6
h1 7 15 7
nc 8 16 8

Associations

MDA : 1/1
MDA : 1/2
Port-id : 1/1/2
Port-id : 1/1/3
Port-id : 1/1/4
Port-id : 1/1/5
Port-id : 1/1/6
Port-id : 1/1/7
Port-id : 1/1/8
=====
*A:PE#

```



## 6 Service Ingress and Egress QoS Policies

### 6.1 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.

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. Other QoS policies must be explicitly associated.

For information about the tasks and commands necessary to access the CLI and to configure and maintain routers, refer to the CLI Usage chapter in the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Basic System Configuration Guide*.

---

## 6.2 Basic Configurations

A basic service egress QoS policy must have the following:

- a unique service egress QoS policy ID
- a QoS policy scope of template or exclusive
- at least one defined default queue

A basic service ingress QoS policy must have the following:

- a unique service ingress QoS policy ID
- a QoS policy scope of template or exclusive
- at least one default unicast forwarding class queue
- at least one multipoint forwarding class queue

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.

## 6.3 Service Ingress QoS Policy

To create a service ingress policy, define the following:

- A policy ID value. The system will not dynamically assign a value.
- A description. The description provides a brief overview of policy features.
- 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.
- A default priority for all packets received on an ingress SAP using this policy.
- Mappings from incoming packet contents to a forwarding class, then separately, from the forwarding class to queue.
- 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**.
- A precedence value for the forwarding class or enqueueing priority when a packet is marked with an IP precedence value.
- 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, 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#
```

### 6.3.1 Service Ingress QoS Queue

To create a service ingress queue, 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#

```

## 6.3.2 Ingress Percent-Rate Support

The **percent-rate** command is supported in a SAP ingress QoS policy for **pir** and **cir** parameters for both queues and policers, with the **fir** parameter supported only for queues on FP4 hardware that is ignored when the related policy is applied to FP2- or FP3-based hardware. For **pir**, the range is 0.01 to 100.00, and for **cir** and **fir**, the range is 0.00 to 100.00.

For queues, when the queue rate is **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. If there is no parent scheduler rate, or its rate is **max**, the **port-limit** is used. When the **port-limit** is used, the **percent-rate** is relative to the rate of the port (including the ingress-rate setting) to which the queue is attached. The **port-limit** is the default.

For policers, the percent-rate rate is always relative to the immediate parent root policer/arbitrator rate or the FP capacity.

The following shows a SAP ingress QoS policy configuration:

```

*A:PE>config>qos>sap-ingress# queue 1 percent-rate
- no percent-rate
- percent-rate <pir-percent> [cir <cir-percent>] [fir <fir-percent>] [port-
limit | local-limit] [fir <fir-percent>]
- percent-rate <pir-percent> police [port-limit | local-limit]

<pir-percent> : [0.01..100.00]
<cir-percent> : [0.00..100.00]
<fir-percent> : [0.00..100.00]
<police> : keyword
<port-limit | local-*> : keyword

*A:PE>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]
```

```
*A:PE>config>qos>sap-ingress#
```

### 6.3.3 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#

```

## 6.3.4 Ingress IP Match Criteria

When specifying SAP ingress match criteria, only one match criteria type (IP/IPv6 or MAC) 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#

```

## 6.3.5 Ingress IPv6 Match Criteria

When specifying SAP ingress match criteria, only one match criteria type (IP/IPv6 or MAC) can be configured in the SAP ingress QoS policy. This feature applies only to the 7750 SR and 7950 XRS.

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 2001:db8:1000::/64
dst-ip 2001:db8:2000::/64
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#
```

### 6.3.6 Ingress MAC Match Criteria

Both IP/IPv6 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 enqueueing 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#
```

### 6.3.6.1 VLAN ID Matching

On ingress, VLAN ID matching may be used to set QoS on SAP ingress. The matching rules are the same as for VID filter (See “VID filters” in the Filter Policies section of the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*) 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 follows) could be constructed, 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
```

### 6.3.7 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 by configuring a policer in the action statement: `ip-criteria`, `ipv6-criteria`, 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.

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 shown in [Figure 12](#)) and directed to policer 1 while the remainder of the customer's traffic is processed using ingress queue 1.

**Figure 12** 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
```

## 6.3.8 Virtual Network Identifier (VNI) Classification

Virtual Network Identifier (VNI) classification is supported for VXLAN and VXLAN GPE traffic within a SAP ingress QoS policy. This classification is configured in the **ip-criteria** and **ipv6-criteria** contexts with **type vxlan-vni** (changed from the default **type normal**). The matching entry must be created with **match protocol udp** for IPv4 or **match next-header udp** for IPv6, and uses the **vxlan-vni** parameter within the **match** statement to match on a single VNI or a range of VNIs.

An alternative to SAP ingress VNI classification is queue group redirection, which provides a more advanced mechanism where individual VNIs can use their own access ingress and egress queue group instances for QoS control (see [VXLAN VNI Queue Group Redirection](#)).

The type cannot be changed when **ip-criteria** or **ipv6-criteria** entries are configured. If there are no **ip-criteria** or **ipv6-criteria** entries configured, the type can be changed from **vxlan-vni** to **normal**. The type can only be changed from **normal** to **vxlan-vni** if there are no **ip-criteria** or **ipv6-criteria** entries configured and if the SAP ingress QoS policy has not been applied to any object.

The following is an example where traffic received with a VNI of 1 is sent to policer 1 and VNIs 2 to 10 are sent to policer 2:

```
sap-ingress 10 create
 queue 1 create
 exit
 queue 11 multipoint create
 exit
 policer 1 create
 exit
 policer 2 create
 exit
 ip-criteria
 type vxlan-vni
 entry 10 create
 match protocol udp
 vxlan-vni eq 1
 exit
 action policer 1
 exit
 entry 20 create
 match protocol udp
 vxlan-vni range 2 10
 exit
 action policer 2
 exit
 exit
exit
```

Ingress VNI classification is applicable to all Ethernet SAPs, except for PW-SAPs, B-VPLS SAPs, and CCAG SAPs, in any applicable service.

The following restrictions also apply:

- Source and destination port matching on a SAP, on which a SAP ingress QoS policy is applied that has **ip-criteria** or **ipv6-criteria** statements with **type vxlan-vni**, is not available for:

- IPv4 QoS classification for VXLAN or VXLAN GPE traffic
- IPv6 QoS and filter classification for VXLAN or VXLAN GPE traffic

If the criteria type is set to **vxlan-vni** and if source or destination port matching entries are configured in an IPv4 or IPv6 SAP ingress QoS policy or in an IPv6 filter policy, any VXLAN or VXLAN GPE ingress traffic will not match these entries on the SAP on which the SAP ingress QoS policy is applied.

- The simultaneous configuration on a SAP of a QoS policy containing an **ip-criteria** entry with **type vxlan-vni** and of a MAC filter with **type vid** is not supported, and vice versa.

This is only applicable to Epipe and VPLS services.

- If a SAP ingress QoS policy that has **ip-criteria** or **ipv6-criteria** statements with **type vxlan-vni** is applied to a SAP, any **ip-criteria** or **ipv6-criteria** entry match **vxlan-vni** statements will not match:

- IPv4 packets containing options
- IPv6 packets containing extension headers
- Ingress 802.1ah PBB frames
- IPv6 over PPPoE traffic received with more than one VLAN tag
- Non-first fragments of an IPv4 or IPv6 fragmented packet

- The configuration of a SAP ingress QoS policy containing **ip-criteria** or **ipv6-criteria** entry match **vxlan-vni** statements is not supported within an SLA profile.

### 6.3.9 FC Mapping Based on EXP Bits

Use the **lsp-exp** command to set the 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 hsmcda-counter-override 1
lsp-exp 3 fc "ef" priority high hsmcda-counter-override 2

```

## 6.3.10 Storing Match Criteria Entries

FP2-, FP3-, and FP4-based cards store QoS policy match-criteria entries in dedicated memory banks in hardware also referred to as CAM tables:

- IP/MAC ingress
- IP/MAC egress
- IPv6 ingress
- IPv6 egress

### 6.3.10.1 FP4-based Cards

To optimize both scale and performance, policy match-criteria entries configured by the operator are compressed by each FP4 line card prior to being installed in hardware.

This compression can result, in an unexpected scenario typically only achieved in a lab environment, in an overload condition where entries for a line card QoS policy on a line card FP are not programmed. This overload condition can occur when applying a QoS policy for the first time on a line card FP or when adding entries to a QoS policy.

#### Applying a QoS Policy

A policy is installed for the first time on a line card FP if no router interface, service interface, SAP, spoke SDP, mesh SDP or ESM subscriber host was using the policy on this FP.

A policy installed for the first time on a line card FP can lead to a compression failure resulting in an overload condition for this policy on this FP. In this case, none of the entries for the affected QoS match-criteria policy are programmed, and the QoS policy queuing, FC mapping, and FC marking are unaffected.

### **Adding QoS Match-Criteria Entries**

Adding an additional entry to a QoS policy can lead to a compression failure resulting in an overload condition.

In this case, the newly added entry is not programmed on the affected FP. Additional entries added to the same policy after the first overload condition are also not programmed on the affected FP as the system attempts to install all outstanding additions in order.

The CPM QoS management task controls the maximum number of match criteria entries per FP. If the operator attempts to go over the scaling limit, the system will return an interactive error message.

**Note:** A trap is raised if a policy is in overload, there is no interactive error message.

### **Removing QoS Match-Criteria Entries**

Removing match-criteria entries from a QoS policy is always successful.

### **Resolving Overload**

The overload condition should be resolved by the network operator before adding new entries in the affected policy.

To identify the affected policy, the system logs the overload event providing slot number, FP number, and impacted memory bank. Based on this information, the **tools>dump>qos>match-criteria-overload** command allows the operator to identify the affected policy and policy entries in the system.

To resolve the overload condition, the network operator can remove the newly added entries from the affected policy or assign a different policy.



## 6.4 Service Egress QoS Policy

To create a service egress policy, define the following:

- A new policy ID value. The system will not dynamically assign a value.
- 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.
- A description. The description provides a brief overview of policy features.

### 6.4.1 Service Egress QoS Queue

To create a service egress QoS queue, define the following:

- 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. 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#
```

## 6.4.2 Egress Percent-Rate Support

The **percent-rate** command is supported in a SAP egress QoS policy for **pir** and **cir** parameters for both queues and policers. For **pir**, the range is 0.01 to 100.00, and for **cir**, the range is 0.00 to 100.00.

For queues, when the queue rate is **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. If there is no parent scheduler rate or **agg-rate rate**, or those rates are **max**, the **port-limit** is used. When the **port-limit** is used, the **percent-rate** is relative to the rate of the port (including the egress-rate setting) to which the queue is attached. The **port-limit** is the default.

For policers, the **percent-rate** rate is always relative to the immediate parent root policer/arbitrator rate or the FP capacity.

The following shows a SAP egress QoS policy configuration:

```
*B:PE>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:PE>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]

*B:PE>config>qos>sap-egress#
```

## 6.4.3 Dynamic MBS for Egress Queue Group Queues

Dynamic MBS is used to constrain the maximum delay experienced by the traffic forwarded through an egress queue group queue when the operational PIR of the queue is modified as part of the HQoS algorithm.

The approximate maximum delay of traffic through a queue due to the length of the queue, when the queue is not using HQoS, is relative to its administrative PIR and can be approximated as  $(\text{MBS}[\text{KB}] \times 8) / \text{PIR}[\text{kb/s}]$  in seconds. A queue's PIR is set to **max**, and its administrative PIR is set to the rate of the port to which the queue is attached.

When using HQoS, the PIR is modified by the HQoS algorithm to give an operational PIR that is equal to or lower than the administrative PIR. As the operational PIR changes, the delay through the queue can also change if the length of the queue is fixed. Reducing the operational PIR could increase the delay, while increasing the operational PIR could reduce the delay. Enabling dynamic MBS on a queue allows the system to change the administrative MBS of the queue in a ratio of operational PIR to administrative PIR, giving an operational MBS, which aims to maintain the maximum queue delay. A queue's drop tails and WRED slope parameters are defined as percentages of the MBS and are, therefore, adjusted accordingly.

When any of the queue parameters are reduced, packets that are already in the queue will not be affected and will be forwarded. Reducing these parameters will constrain the latency for newly arriving packets, but those packets already in the queue before the new parameter values were set will be forwarded with the delay associated with the actual queue depth when the packet was enqueued (based on the previous parameter values).

The configured CBS is used as a minimum operational MBS. The maximum MBS is capped by the maximum administrative MBS (1 GB).

If the operational MBS changes such that its value is similar or equal to the configured CBS, the system increases the CBS to ensure that buffers can be requested from the correct portion of the buffer pool (shared or reserved). This operation is automatic and the CBS reverts to its configured value if the MBS is increased sufficiently. The automatic increase in the CBS could, however, cause the **resv-cbs** red or amber alarms to be raised if the increase in the related queues' CBS results in the total CBS assigned (but not necessarily used) matching or exceeding the **resv-cbs** red and amber thresholds.

If a LAG is used together with **pool-per-queue**, the related hardware queues exist in their own pool in the egress WRED megapool on a given FP and the operational MBS is used to size the shared part of the pool with the sum of the CBS defining the reserved part of the pool.

Dynamic MBS is supported for both native FP and **pool-per-queue** queues within an egress queue group template, which can be applied to access or network Ethernet ports and used for egress network interface traffic, egress SAP traffic, and subscriber egress policed traffic.

The configuration of dynamic MBS and queue depth monitoring are mutually exclusive.

Dynamic MBS is configured as follows:

```
configure
 qos
 queue-group-templates
 egress
 queue-group <queue-group-name> create
 queue <queue-id>
```

The operational MBS can be shown using the **show pools** and **show qos scheduler-hierarchy** commands.

The following example shows the use of dynamic MBS. A queue group template is applied to port 5/1/1 configured with multiple queues using HQoS, one of which has the following parameters:

```
queue-group "qg1" create
 queue 1 best-effort create
 parent "s1"
 rate 50000
 mbs 1000 kilobytes
 dynamic-mbs
 exit
```

Without any traffic in the other queues constraining the operational PIR on this queue, the MBS used is the administrative MBS.

```
*B:PE# show pools access-egress 5/1/1 queue-group "qg1" instance 1
...
=====
Queue : accQGrp->qg1:1(5/1/1)->1
=====
FC Map : not-applicable
Tap : not-applicable FP : not-applicable
Admin PIR : 50000 Oper PIR : 50000
Admin CIR : 0 Oper CIR : 0
Admin MBS : 1008 KB Oper MBS : 1008 KB
High-Plus Drop Tail : 1008 KB High Drop Tail : 1008 KB
Low Drop Tail : 888 KB Exceed Drop Tail : 792 KB
CBS : 0 KB Depth : 0
Slope : not-applicable
```

If traffic is sent to the other queues in the queue group such that the operational PIR of queue 1 is reduced to 25 Mb/s, the show output changes to:

```
*B:PE# show pools 5/1/1 access-egress queue-group "qg1" instance 1
...
=====
```

```
Queue : accQGrp->qg1:1(5/1/1)->1
=====
FC Map : not-applicable
Tap : not-applicable FP : not-applicable
Admin PIR : 50000 Oper PIR : 25000
Admin CIR : 0 Oper CIR : 0
Admin MBS : 1008 KB Oper MBS : 504 KB
High-Plus Drop Tail : 1008 KB High Drop Tail : 1008 KB
Low Drop Tail : 948 KB Exceed Drop Tail : 900 KB
CBS : 0 KB Depth : 0
Slope : not-applicable
```

The output shows that the operational MBS is now 50% of the administrative MBS and that the queue's drop tails have changed accordingly.

### 6.4.3.1 Queue Length as a Delay Value

The length of a queue within an egress queue group template can be configured as a target queue delay, in ms, rather than an absolute byte/kbytes value. The queue MBS is calculated from the queue delay and the administrative PIR, with the MBS [kB] being approximately the value of  $((\text{queue delay}[\text{ms}]/1000) \times (\text{PIR}[\text{kb/s}] / 8))$ .

The queue delay is configured as follows:

```
configure
 qos
 queue-group-templates
 egress
 queue-group <queue-group-name> create
 queue <queue-id>
 queue-delay <ms>
```

The **queue-delay** command and the **mbs** command are mutually exclusive. To change between the **mbs** and **queue-delay** parameters, the current parameter must be removed before adding the new parameter; that is, changing from **mbs** to **queue-delay** requires a **no mbs** before the **queue-delay** is configured and changing from **queue-delay** to **mbs** requires a **no queue-delay** before the **mbs** is configured. If **queue-delay** is configured for an egress queue group queue, it is not possible to override the MBS for that queue.

The following is an example of configuring queue delay and the resulting MBS:

```
*A:PE# configure qos queue-group-templates egress queue-group "qg2"
*A:PE>cfg>qos>qgrps>egr>qgrp# info

queue 1 best-effort create
 rate 50000
 queue-delay 100
```

```

 exit

*A:PE>cfg>qos>qgrps>egr>qgrp# exit all
*A:PE# show pools access-egress 5/1/1 queue-group "qg2" instance 1
...
=====
Queue : accQGrp->qg2:1(5/1/1)->1
=====
FC Map : not-applicable
Tap : not-applicable FP : not-applicable
Admin PIR : 50000 Oper PIR : 50000
Admin CIR : 0 Oper CIR : 0
Admin MBS : 612 KB Oper MBS : 612 KB
High-Plus Drop Tail : 612 KB High Drop Tail : 612 KB
Low Drop Tail : 540 KB Exceed Drop Tail : 480 KB
CBS : 0 KB Depth : 0
Slope : not-applicable

```

## 6.4.4 Egress SAP FC and FP 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 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.

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.

Dot1p, 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. Dot1p, 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 Dot1p, 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.

## 6.4.5 Egress IP Match Criteria

The following displays a SAP egress QoS policy IP criteria configuration:

```
*A:PE>config>qos# info

#-----
echo "QoS Policy Configuration"
#-----
 sap-egress 10 name "10" create
 queue 1 create
 exit
 ip-criteria
 entry 10 create
 match
 src-ip 192.168.0.0/16
 exit
 action fc "af"
 exit
 entry 20 create
 match
 dst-ip 10.0.0.0/8
 exit
 action fc "be"
 exit
 exit
 exit

*A:PE>config>qos#
```

## 6.4.6 Egress IPv6 Match Criteria

The following displays a SAP egress QoS policy IPv6 criteria configuration:

```
*A:PE>config>qos# info

#-----
echo "QoS Policy Configuration"
#-----
 sap-egress 10 name "10" create
 queue 1 create
 exit
 ipv6-criteria
 entry 10 create
 match
 dst-ip 2001:db8:1000::/64
 exit
 action profile in fc "ef"
 exit
 exit
 exit

```

---

```
*A:PE>config>qos#
```

## 6.4.7 Egress 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 egress by configuring a policer in the **action** statement within an **ip-criteria** or **ipv6-criteria** statement.

The policed traffic by default exits through one of the following methods:

- A queue in the **policer-output-queues queue** group that is automatically created on an access or hybrid port with the queue used that was chosen by the forwarding class definition in that queue group. If the forwarding class is modified in the **action** statement, the new forwarding class selects the queue to be used.
- A specific queue in a user-configured queue group. For SAP egress, this requires the use of the **port-redirect-queue-group queue** parameter in the criteria **action** statement with the queue group name being specified when the egress QoS policy is applied to the SAP. For subscribers, the queue group to be used is selected using the inter-dest-id associated with the subscriber and configured as the **host-match dest** under the port access queue group configuration.
- A SAP queue configured within the SAP egress QoS policy.
- The queue to which the forwarding class for the traffic is mapped. This could be a queue group, SAP, or subscriber queue. This requires the use of the **use-fc-mapped-queue** parameter in the criteria **action** statement. If the forwarding class is modified in the **action** statement, new forwarding class selects the queue to be used.

The number of configuration combinations of a policer and one of the preceding methods is capped at 63 within a given SAP egress QoS policy. For two or more definitions to be counted as a single combination, their action statement must have the same policer ID, the same queue ID (if specified in either statement), the same **port-redirect-queue-group** (if specified in either statement), and the parameter **use-fc-mapped-queues** (if specified in either statement).



The forwarding class and profile used are irrelevant when considering the number of combinations. For example, it is possible to configure 32 policers with traffic exiting queue 1, but then, only 31 of the same policers are exiting queue 2; this would use all 63 combinations. A resource is also allocated per FP where each combination configured corresponds to an egress bypass entry used in the FP per sap-instance or per subscriber-sap-sla instance that uses the egress qos policy. The number of egress bypass entries available on an FP, together with the number allocated and the number free, can be seen using the following tools command.

```
*A:PE>tools>dump>resource-usage# card 1 fp 1
=====
Resource Usage Information for Card Slot #1 FP #1
=====

Total Allocated Free


...
Egress Qos Bypass | 262143 0 262143
...

```

This is supported on all hardware systems, excluding when a HSMDA is used. QPPB processing takes precedence over this feature.

This could be used, for example, when it is required that egress traffic with a DSCP value EF is to be policed instead of shaped in a queue on a given SAP. The traffic could be classified based on its DSCP value and directed to **policer 1** while the remainder of the customer's traffic is processed using egress **queue 1**. This is shown in [Figure 13](#).

**Figure 13 Egress SAP**



The configuration would be as follows:

```
sap-egress 10 create
queue 1 create
exit
policer 1 create
exit
ip-criteria
```

```
entry 10 create
 match
 dscp ef
 exit
 action policer 1
exit
exit
exit
```

## 6.4.8 Storing Match Criteria Entries

FP2-, FP3-, and FP4-based cards store QoS policy match-criteria entries in dedicated memory banks in hardware, also referred to as CAM tables:

- IP/MAC ingress
- IP/MAC egress
- IPv6 ingress
- IPv6 egress

### 6.4.8.1 FP4-based Cards

To optimize both scale and performance, policy match-criteria entries configured by the operator are compressed by each FP4 line card prior to being installed in hardware.

This compression can result, in an unexpected scenario typically only achieved in a lab environment, in an overload condition where entries for a line card QoS policy on a line card FP are not programmed. This overload condition can occur when applying a QoS policy for the first time on a line card FP or when adding entries to a QoS policy.

#### Applying a QoS Policy

A policy is installed for the first time on a line card FP if no router interface, service interface, SAP, spoke SDP, mesh SDP or ESM subscriber host was using the policy on this FP.

A policy installed for the first time on a line card FP can lead to a compression failure resulting in an overload condition for this policy on this FP. In this case, none of the entries for the affected QoS match-criteria policy are programmed, and the QoS policy queuing, FC mapping, and FC marking are unaffected.

### Adding QoS Match-Criteria Entries

Adding an additional entry to a QoS policy can lead to a compression failure resulting in an overload condition.

In this case, the newly added entry is not programmed on the affected FP. Additional entries added to the same policy after the first overload condition are also not programmed on the affected FP as the system attempts to install all outstanding additions in order.

The CPM QoS management task controls the maximum number of match criteria entries per FP. If the operator attempts to go over the scaling limit, the system will return an interactive error message.

**Note:** A trap is raised if a policy is in overload, there is no interactive error message.

### Removing QoS Match-Criteria Entries

Removing match-criteria entries from a QoS policy is always successful.

### Resolving Overload

The overload condition should be resolved by the network operator before adding new entries in the affected policy.

To identify the affected policy, the system logs the overload event providing slot number, FP number, and impacted memory bank. Based on this information, the **tools>dump>qos>match-criteria-overload** command allows the operator to identify the affected policy and policy entries in the system.

To resolve the overload condition, the network operator can remove the newly added entries from the affected policy or assign a different policy.

## 6.4.9 Dot1p Egress Remarking

Dot1p remarking can be performed on egress for all services and with respect to the profile of the packet and the VLAN tag.

The following commands can be used to remark the dot1p values at a SAP egress:

**CLI Syntax:**

```
configure qos sap-egress policy-id create
fc fc-name create
dot1p {dot1p-value | in-profile dot1p-value out-
profile dot1p-value [exceed-profile dot1p-value]}
```

```
dot1p-inner {dot1p-value | in-profile dot1p-value
out-profile dot1p-value}
dot1p-outer {dot1p-value | in-profile dot1p-value
out-profile dot1p-value [exceed-profile dot1p-
value]}
exit
```

All inplus-profile traffic is marked with the same value as in-profile traffic.

The precedence of these commands is summarized as follows, 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, 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. If qinq-mark-top-only is omitted, both the inner and outer VLAN tags are remarked.

Remarking the inner dot1p is not supported based on the profile result of egress policing.

The egress remarking occurs after any egress classification.

### 6.4.9.1 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 determined using existing QoS fields; for example, the three MPLS EXP and Ethernet dot1p bits. Using certain combinations of these bits to indicate both forwarding class (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 (priorities) that could be indicated using the three 802.1p bits. 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 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
```

By default, the DE bit is set to 0 for inplus-profile and in-profile traffic and 1 for out-of-profile and exceed-profile traffic, unless explicitly forced.

The precedence of these commands is summarized from highest to lowest precedence, as follows:

- 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, 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. If qinq-mark-top-only is omitted, both the inner and outer VLAN tags are remarked.


Remarking the inner DE bit is not supported based on the profile result of egress policing.

The egress remarking occurs after any egress classification.

6.4.9.1.1 DEI in IEEE 802.1ad

The *IEEE 802.1ad-2005 standard* allows drop eligibility to be conveyed separately from priority in-service VLAN TAGs (S-TAGs). The S-TAG has a new format where the priority and discard eligibility parameters are conveyed in the 3-bit priority code point (PCP) field and in the DE bit, respectively (see [Figure 14](#)).

**Figure 14** DE Bit in the 802.1ad S-TAG



The introduction of the DE bit allows the S-TAG to convey eight forwarding classes/ distinct 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. In case of congestion, these packets will be the first ones to be dropped.

6.4.9.1.2 DEI in IEEE 802.1ah

IEEE 802.1ah (PBB) standard provides a dedicated bit for DE indication in both the backbone VLAN ID (BVID) and the ITAG.

The BVID is a regular 802.1ad S-TAG. 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](#).

### 6.4.9.1.3 IEEE 802.1ad Use Case

**Figure 15** shows 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 the following 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 S-TAG.


**Figure 15 DE Aware 802.1ad Access Network**



Fig\_26

**Figure 16** shows an example of the QoS processing of the DEI processing steps for the IEEE 802.1ad use case for both ingress and egress directions (from a PE1 SAP perspective).

**Figure 16 DEI Processing Ingress into the PE1 SAP**



Fig\_28

The following steps related to DEI are involved in the QoS processing as the packet moves from left to right:

1. The QinQ access device sets the DE bit from the S-TAG based on the QoS classification or on the results of the metering/policing for the corresponding customer UNI.


2. The SAP on PE1 may use the DE bit from the customer S-TAG to classify the frames as in/out-of-profile. Color aware policing/metering can generate additional out-of-profile packets as the result of packet flow surpassing the CIR.
3. 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).
4. As the packet arrives at PE2, on ingress into the related SDP, the DE indication is used to classify the packets into an internal PHB.
5. On 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.

#### 6.4.9.1.4 IEEE 802.1ah Use Case

Figure 17 shows an example of a PBB topology where the DE feature can be used. The processing requirements highlighted in the 802.1ad use case apply to the 802.1ah BVID, format and etype, these being identical with the 802.1ad S-TAG. 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/S-TAG; for example, the DE bit from the BVID header represents the QoS associated with the “Ethernet Tunnel” while the DE bit from the ITAG represents the service QoS.

**Figure 17 DE Aware PBB Topology**



Fig\_27

In this example, the BVID is not used for a part of the network, leaving the I-DEI bit from the ITAG as the only option for a dedicated DE field. If both are included, the QoS information from the BVID is to be used.



## 6.4.10 DSCP and IP Precedence Egress Remarking

DSCP and IP precedence remarking can be performed on egress for all services and with respect to the profile of the packet.

Use the following CLI syntax to remark the DSCP and IP precedence values at a SAP egress:

**CLI Syntax:**

```
configure qos sap-egress policy-id create
 fc fc-name create
 dscp {dscp dscp-name | in-profile dscp-name
 out-profile dscp-name [exceed-profile dscp-
 name]}
 prec {ip-prec-value | in-profile ip-prec-value
 out-profile ip-prec-value} [exceed-profile
 ip-prec-value]}
 exit
exit
```

All inplus-profile traffic is marked with the same value as in-profile traffic.

Remarking the DSCP and IP precedence based on the profile result of egress policing must be enabled under the related policer configuration, as follows:

**CLI Syntax:**

```
sap-egress <policy-id> create
 policer <policer-id> create
 enable-dscp-prec-remarking
 exit
exit
```

## 6.4.11 Queue Depth Monitoring

Queue depth monitoring gives more visibility to the user of the queue depths being experienced on a set of queues when the traffic is bursty. The instantaneous depth of a queue can be seen using the **show pools** command, whereas queue depth monitoring shows the variation in queue depth over a period of time. It is applicable to SAP ingress unicast and multipoint queues and SAP egress queues, and for ingress and egress access and network queue group queues used by any service or network interfaces. The monitoring uses a polling mechanism by the line card CPU. Consequently, the results provided are statistical.

An override (**monitor-depth**) is used to enable queue depth monitoring, which is configured under the SAP or queue group queue-overrides. There are show and clear commands, using the **queue-depth** parameter, for both service SAPs and port queue groups with associated MIB variables.

The following configuration shows an example of enabling the monitoring of the depth of queue 1 on an Epipe SAP.

```
epipe 1 customer 1 create
 sap 1/2/1 create
 egress
 qos 10
 queue-override
 queue 1 create
 monitor-depth
 exit
 exit
 exit
exit
```

The queue depth can then be shown as follows:

```
*A:PE-1# show service id 1 sap 1/2/1 queue-depth

=====
Queue Depth Information (Ingress SAP)
=====
No Matching Entries
=====

=====
Queue Depth Information (Egress SAP)
=====

Name : 1->1/2/1->1
MBS : Def

Queue Depths (percentage)

0%-10% 11%-20% 21%-30% 31%-40% 41%-50% 51%-60% 61%-70% 71%-80% 81%-90% 91%-100%

68.21 3.64 3.43 3.47 3.86 3.22 3.86 2.87 3.78 3.66

Average Elapsed Time : 0d 00:11:48
Wghtd Avg Polling Interval: 99 ms

=====
*A:PE-1#
```

The output shows the percentage of polls for each 10% range of queue depth. The output includes the name of the queue, its MBS configuration, the average elapsed time over which the depth was monitored (this is the elapsed time since the start of monitoring or the last clear), and the weighted average polling interval.

For example, in the preceding output, the queue depth was in the range of 51% to 60% for 3.22% of the polls, the polling was performed over an elapsed time of 11 minutes and 48 seconds, and the average polling interval was 99 ms.

The monitoring is performed on the hardware queues corresponding to the configured queue. It is possible that the set of related hardware queues for a given configured queue changes over time; for example, when LAG ports are added or removed resulting in monitored hardware queues being added or removed. If the set of hardware queues for the configured queue changes, the system will only report occupancy information of all currently instantiated hardware queues; no attempt is made to keep historical occupancy information.

The average polling interval is weighted based on the elapsed monitoring time of the individual hardware queues corresponding to the configured queue, and the elapsed monitoring time is averaged over the same set of hardware queues.

There is no specific limit on the number of queues that can be monitored, but the amount of each line card CPU's resources allocated to the monitoring is bounded. Consequently, the average polling interval will increase as more queues are monitored on the line card.

If the MBS of a queue is modified, the occupancy information is cleared and the elapsed timers reset to zero. Issuing a **clear card** command will also clear this information. Packet drops caused at the pool level, rather than at the queue level, would result in lower queue depths being reported.

## 6.5 Service Management Tasks

This section discusses service ingress and egress service management tasks.

### 6.5.1 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 *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide* for information about configuring service parameters on the 7750 SR and 7450 ESS.

#### 6.5.1.1 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#
```

### 6.5.1.2 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
 exit
no shutdown
exit

```

### 6.5.1.3 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#
```

### 6.5.1.4 VPRN

The following output displays a VPRN service configuration for the 7750 SR and 7950 XRS.

```
A:ALA-7>config>service# info

...
 vprn 1 customer 1 create
 ecmp 8
 autonomous-system 10000
 route-distinguisher 10001:1
 auto-bind-tunnel
 resolution-filter
 resolution-filter ldp
 vrf-target target:10001:1
 interface "to-cel" create
 address 10.1.0.1/24
 sap 1/1/10:1 create
 ingress
 qos 100
 exit
 egress
 qos 105
 exit
 exit
 exit
 no shutdown
 exit
...

A:ALA-7>config>service#
```

## 6.5.2 Editing QoS Policies

QoS existing policies and entries can be edited. 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, then write over the original policy.

## 6.5.3 Copying and Overwriting QoS Policies

Existing service egress or ingress policy can be copied, renamed 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#
```

## 6.5.4 Deleting QoS Policies

Every service SAP is associated, by default, with the appropriate egress or ingress policy (policy-id 1). The default policy can be replaced with a customer-configured policy, but cannot entirely remove the policy from the SAP configuration. When a non-default service egress or ingress policy is removed, the association reverts to the default policy-id 1.

---

A QoS policy cannot be deleted until it is removed from all SAPs where it is 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#
```

## 6.5.5 Removing 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`



## 6.6 Service Ingress and Egress QoS Policy Command Reference

### 6.6.1 Command Hierarchies

- [Service Ingress QoS Policy Commands](#)
- [Service Egress QoS Policy Commands](#)
- [Operational Commands](#)
- [Show Commands](#)
- [Debug Commands](#)

#### 6.6.1.1 Service Ingress QoS Policy Commands

```

config
— qos
— sap-ingress {policy-id | policy-name} [create] [name name]
— no sap-ingress {policy-id | policy-name}
— default-fc fc
— 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 [dscp-name] [fc fc-name] [priority {low | high}]
— no dscp dscp-name [dscp-name]
— dynamic-policer
— cbs {size [bytes | kilobytes] | default}
— no cbs
— mbs {size [bytes | kilobytes] | default}
— no mbs
— packet-byte-offset {add add-bytes | subtract sub-bytes}
— no packet-byte-offset
— parent arbiter-name [weight weight-level] [level level]
— no parent
— range start-entry policer-id count count
— no range
— 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

```

- 
- **fc** *fc-name* [create]
  - **no fc** *fc-name*
    - **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* | **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* | **prec** *ip-prec-value*}
    - **no out-remark**
    - **policer** *policer-id* [fp-redirect-group]
    - **no policer**
    - **profile** {in | out}
    - **no profile**
    - **queue** *queue-id* [group *queue-group-name*]
    - **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**
    - **entry** *entry-id* [create]
    - **no entry** *entry-id*
      - **action** [**fc** *fc-name*] [priority {high | low}] [**policer** *policer-id*]
      - **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*] | **ip-prefix-list** *prefix-list-name*}
        - **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/mask* | *ip-address* [*netmask*] | **ip-prefix-list** *prefix-list-name*}
        - **no src-ip**
        - **src-port** {lt | gt | eq} *src-port-number*
        - **src-port range** *start end*
        - **no src-port**
        - **vxlan-vni** [eq *vxlan-vni-id* | **range** *start end*]
        - **no vxlan-vni**

```

— renum old-entry-id new-entry-id
— type filter-type
— no type
— [no] ipv6-criteria
— entry entry-id [create]
— no entry entry-id
— action [fc fc] [priority {low | high}] [policer policer-id]
— 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 | ipv6-address ipv6-address-mask}
— no dst-ip
— dst-port {lt | gt | eq} dst-port-number
— dst-port range start end
— no dst-port
— fragment {true | false | first-only | non-first-only}
— no fragment
— src-ip {ipv6-address/prefix-length | ipv6-address ipv6-address-mask}
— no src-ip
— src-port {lt | gt | eq} src-port-number
— src-port range start end
— no src-port
— vxlan-vni [eq vni-id | range start end]
— no vxlan-vni
— renum old-entry-id new-entry-id
— type filter-type
— no type
— isp-exp isp-exp-value [fc fc-name] [priority {low | high}]
— no isp-exp isp-exp-value
— [no] mac-criteria
— entry entry-id [create]
— no entry entry-id
— action [fc fc] [priority {low | high}] [policer policer-id]
— no action
— description description-string
— no description
— match [frame-type {802dot3 | 802dot2-llc | 802dot2-snap | ethernet-II | atm}]
— no match
— atm-vci vci-id
— 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**
- **inner-tag** *value* [*vid-mask*]
- **no inner-tag**
- **outer-tag** *value* [*vid-mask*]
- **no outer-tag**
- **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*
- **type** *filter-type*
- **no type**
- **policer** *policer-id* [*create*]
- **no policer** *policer-id*
  - **adaptation-rule** [*pir adaptation-rule*] [*cir adaptation-rule*]
  - **no adaptation-rule**
  - **adv-config-policy** *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 add-bytes* | *subtract sub-bytes*}
  - **no packet-byte-offset**
  - **parent** *arbiter-name* [*weight weight-level*] [*level level*]
  - **no parent**
  - **percent-rate** *pir-percent* [*cir cir-percent*]
  - **no percent-rate**
  - [**no**] **profile-capped**
  - **rate** {*max* | *pir-rate*} [*cir* {*max* | *cir-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**
- **prec** *ip-prec-value* [*fc fc-name*] [*priority* {*low* | *high*}]
- **no prec** *ip-prec-value*
- **queue** *queue-id* [*multipoint*] [*queue-type*] [*queue-mode*] [*create*]
- **no queue** *queue-id*
  - **adaptation-rule** [*pir adaptation-rule*] [*cir adaptation-rule*] [*fir* {*max* | *min* | *closest*}]
  - **no adaptation-rule**
  - **adv-config-policy** *policy-name*
  - **no adv-config-policy**
  - **burst-limit** *size* [*bytes* | *kilobytes*]
  - **no burst-limit**

- **cbs** *size-in-kbytes*
- **no cbs**
- **[no] cir-non-profiling**
- **drop-tail**
  - **low**
    - **percent-reduction-from-mbs** *percent*
    - **no percent-reduction-from-mbs**
- **mbs** *size* [*bytes* | *kilobytes*]
- **no mbs**
- **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
- **no packet-byte-offset**
- **parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
- **no parent**
- **percent-rate** *pir-percent* [**cir** *cir-percent*] [**fir** *fir-percent*] [{**port-limit** | **local-limit**}]
- **percent-rate** *pir-percent* **police** [{**port-limit** | **local-limit**}]
- **no percent-rate**
- **pool** *pool-name*
- **no pool**
- **rate** *pir-rate* [**cir** *cir-rate*] [**fir** *fir-rate*]
- **rate** *pir-rate* **police**
- **no rate**
- **scope** {**exclusive** | **template**}
- **no scope**
- **sub-insert-shared-pccrule** **start-entry** *entry-id* **count** *count*
- **no sub-insert-shared-pccrule**

### 6.6.1.2 Service Egress QoS Policy Commands

- ```
config
  qos
    — sap-egress {policy-id | policy-name} [create] [name name]
```
- **no sap-egress** {*policy-id* | *policy-name*}
 - **description** *description-string*
 - **no description**
 - **dot1p** *dot1p-value* [**fc** *fc-name*] [**profile** {**in** | **out** | **use-de** | **exceed** | **inplus**}]
 - **no dot1p** *dot1p-value*
 - **dscp** *dscp-name* [*dscp-name*] [**hsmdda-counter-override** *counter-id*] [**fc** *fc-name*] [**profile** {**in** | **out** | **exceed** | **inplus**}]
 - **no dscp** *dscp-name* [*dscp-name*]
 - **dynamic-policer**
 - **cbs** {*size* [*bytes* | *kilobytes*] | **default**}
 - **no cbs**
 - **mbs** {*size* [*bytes* | *kilobytes*] | **default**}
 - **no mbs**
 - **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
 - **no packet-byte-offset**
 - **parent** *arbiter-name* [**level** *level*] [**weight** *weight-level*]
 - **no parent**
 - **range** **start-entry** *policer-id* **count** *count*

-
- **no range**
 - **stat-mode** {no-stats | minimal | offered-profile-no-cir | offered-profile-cir | offered-total-cir | offered-limited-capped-cir | offered-profile-capped-cir}
 - **no stat-mode**
 - [no] **ethernet-ctag**
 - **fc** *fc-name* [create]
 - **no fc** *fc-name*
 - **de-mark** [force *de-value*]
 - **no de-mark**
 - **de-mark-inner** [force *de-value*]
 - **no de-mark-inner**
 - **de-mark-outer** [force *de-value*]
 - **no de-mark-outer**
 - **dot1p** {*dot1p-value* | in-profile *dot1p-value* out-profile *dot1p-value* [exceed-profile *dot1p-value*]}
 - **no dot1p**
 - **dot1p-inner** {*dot1p-value* | in-profile *dot1p-value* out-profile *dot1p-value*}
 - **no dot1p-inner**
 - **dot1p-outer** {*dot1p-value* | in-profile *dot1p-value* out-profile *dot1p-value* [exceed-profile *dot1p-value*]}
 - **no dot1p-outer**
 - **dscp** {*dscp-name* | in-profile *dscp-name* out-profile *dscp-name* [exceed-profile *dscp-name*]}
 - **no dscp**
 - **hsmdda**
 - **queue** *hsmdda-queue-id* [port-redirect-group-queue]
 - **no queue**
 - **policer** *policer-id* [{[port-redirect-group-queue] [queue *queue-id*] | group *group-name* [instance *instance-id*] [queue *queue-id*]}]
 - **no policer**
 - **prec** {*ip-prec-value* | in-profile *ip-prec-value* out-profile *ip-prec-value* [exceed-profile *ip-prec-value*]}
 - **no prec**
 - **queue** *queue-id* [{group *queue-group-name* [instance *instance-id*] | port-redirect-group-queue}]
 - **no queue**
 - **hsmdda-queues**
 - **low-burst-max-class** *class*
 - **no low-burst-max-class**
 - **packet-byte-offset** {add *add-bytes* | subtract *sub-bytes*}
 - **no packet-byte-offset**
 - [no] **queue** *queue-id*
 - **adaptation-rule** [pir {max | min | closest}]
 - **no adaptation-rule**
 - **burst-limit** *size* [bytes | kilobytes]
 - **no burst-limit**
 - **mbs** {*size* [bytes | kilobytes] | default}
 - **no mbs**
 - **rate** {*pir-rate* | max}
 - **no rate**
 - **slope-policy** *hsmdda-slope-policy-name*
 - **no slope-policy**

-
- **wrr-weight** *weight*
 - **no wrr-weight**
 - **wrr-policy** *wrr-policy-name*
 - **no wrr-policy**
 - [no] **ip-criteria**
 - **entry** *entry-id* [create]
 - **no entry** *entry-id*
 - **action** [hsmda-counter-override *counter-id*] [fc *fc-name*] [profile {in | out | exceed | inplus}] [policer *policer-id*] [port-redirect-group-queue] [queue *queue-id*] [use-fc-mapped-queue]
 - **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*] | *ip-prefix-list* *prefix-list-name*}
 - **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/mask* | *ip-address* [*ipv4-address-mask*] | *ip-prefix-list* *prefix-list-name*}
 - **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**
 - **entry** *entry-id* [create]
 - **no entry** *entry-id*
 - **action** [hsmda-counter-override *counter-id*] [fc *fc-name*] [profile {in | out | exceed | inplus}] [policer *policer-id*] [port-redirect-group-queue] [queue *queue-id*] [use-fc-mapped-queue]
 - **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* | *ipv6-address* *ipv6-address-mask*}
 - **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 | ipv6-address ipv6-address-mask}
— no src-ip
— src-port {lt | gt | eq} src-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 [create]
— no policer policer-id
— adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
— no adaptation-rule
— adv-config-policy policy-name
— no adv-config-policy
— cbs {size [bytes | kilobytes] | default}
— description description-string
— no description
— [no] enable-dscp-prec-remarking
— [no] enable-exceed-pir
— high-prio-only percent-of-mbs
— no high-prio-only
— mbs {size [bytes | kilobytes] | default}
— no mbs
— packet-byte-offset {add add-bytes | subtract sub-bytes}
— no packet-byte-offset
— parent arbiter-name [level level] [weight weight-level]
— no parent
— percent-rate pir-percent [cir cir-percent]
— no percent-rate
— port-parent [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]
— no port-parent
— [no] profile-capped
— [no] profile-out-preserve
— rate {max | pir-rate} [cir {max | cir-rate}]
— scheduler-parent scheduler-name [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]
— no scheduler-parent
— stat-mode {no-stats | minimal | offered-profile-no-cir | offered-profile-cir | offered-total-cir | offered-limited-capped-cir | offered-profile-capped-cir | offered-total-cir-exceed | offered-four-profile-no-cir | offered-total-cir-four-profile}
— no stat-mode
— [no] policers-hqos-manageable
— post-policer-mapping mapping-policy-name
— no post-policer-mapping
— prec ip-prec-value [hsmda-counter-override counter-id] [fc fc-name] [profile {in | out | exceed | inplus}]
— no prec ip-prec-value
— queue queue-id [queue-type] [create]
— 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** {**default** | *size* [**bytes** | **kilobytes**]}
- **no burst-limit**
- **cbs** *size-in-kbytes*
- **no cbs**
- **drop-tail**
 - **exceed**
 - **percent-reduction-from-mbs** *percent*
 - **no percent-reduction-from-mbs**
 - **high**
 - **percent-reduction-from-mbs** *percent*
 - **no percent-reduction-from-mbs**
 - **highplus**
 - **percent-reduction-from-mbs** *percent*
 - **no percent-reduction-from-mbs**
 - **low**
 - **percent-reduction-from-mbs** *percent*
 - **no percent-reduction-from-mbs**
- **mbs** *size* [**bytes** | **kilobytes**]
- **no mbs**
- **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
- **no packet-byte-offset**
- **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**
- **pool** *pool-name*
- **no pool**
- **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**
- **wred-queue** [**policy** *slope-policy-name*] [**mode** *mode*] [**slope-usage** *slope-usage*]
- **no wred-queue**
- **scope** {**exclusive** | **template**}
- **no scope**
- **sub-insert-shared-pccrule** **start-entry** *entry-id* **count** *count*
- **no sub-insert-shared-pccrule**

6.6.1.3 Operational Commands

- ```
config
— qos
— copy hsmda-pool-policy src-name dst-name [overwrite]
```

- **copy** **hsmdda-scheduler-policy** *src-name dst-name* [**overwrite**]
- **copy** **hsmdda-slope-policy** *src-name dst-name* [**overwrite**]
- **copy** **named-pool-policy** *src-name dst-name* [**overwrite**]
- **copy** **sap-egress** *src-pol dst-pol* [**overwrite**]
- **copy** **sap-ingress** *src-pol dst-pol* [**overwrite**]

### 6.6.1.4 Show Commands

- ```
show
  — qos
    — bcg burst-control-group-name [member-queues] [exp-util-bw mega-bits-per-second]
    — bcg burst-control-group-name member-queues [at-risk-only] [exp-util-bw mega-bits-per-second] [port port-id] [sap sap-id] [subscriber sub-ident] [access-queue-grp queue-group-name] [service-id service-id]
    — bcg card-slot-number [ingress | egress]
    — queue queue-group port-id queue-group-name {ingress | egress} [id queue-id] [instance instance-id]
    — queue sap sap-id {ingress | egress} [id queue-id]
    — queue subscriber sub-ident-string [sla-profile sla-profile-name] {ingress | egress} [id queue-id]
    — sap-egress [policy-id] [association | match-criteria | hsmdda | detail]
    — sap-egress summary
    — sap-ingress [policy-id] [association | match-criteria | hsmdda | detail]
    — sap-ingress summary
```

6.6.1.5 Debug Commands

- ```
tools
 — dump
 — qos
 — match-criteria-overload
 — match-criteria-overload card slot-number fp fp-number
 — match-criteria-overload card slot-number fp fp-number network network-policy-id direction direction
 — match-criteria-overload card slot-number fp fp-number sap-egress policy-id
 — match-criteria-overload card slot-number fp fp-number sap-ingress policy-id
```

## 6.6.2 Command Descriptions

- [Generic Commands](#)
- [Service Ingress QoS Policy Commands](#)
- [Service Egress QoS Policy Commands](#)
- [Operational Commands](#)

- [Show Commands](#)
- [Debug Commands](#)

### 6.6.2.1 Generic Commands

#### description

|                    |                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>description</b> <i>description-string</i><br><b>no description</b>                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>qos>sap-egress<br>config>qos>sap-egress>ip-criteria>entry<br>config>qos>sap-egress>ipv6-criteria>entry<br>config>qos>sap-egress>policer<br>config>qos>sap-ingress<br>config>qos>sap-ingress>ip-criteria>entry<br>config>qos>sap-ingress>ipv6-criteria>entry<br>config>qos>sap-ingress>mac-criteria>entry<br>config>qos>sap-ingress>policer    |
| <b>Description</b> | <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The description command associates a text string with a configuration context to help identify the context in the configuration file.</p> <p>The <b>no</b> form of this command removes any description string from the context.</p> |
| <b>Parameters</b>  | <i>description-string</i> — 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.                                                        |

### 6.6.2.2 Service Ingress QoS Policy Commands

#### sap-ingress

|                |                                                                                                                                                                               |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>sap-ingress</b> { <i>policy-id</i>   <i>policy-name</i> } [ <b>create</b> ] [ <b>name</b> <i>name</i> ]<br><b>no sap-ingress</b> { <i>policy-id</i>   <i>policy-name</i> } |
| <b>Context</b> | config>qos                                                                                                                                                                    |

**Description** This command is used to create or edit the ingress policy. The ingress policy defines the SLA enforcement that service packets receive as they ingress a SAP. SLA enforcement is accomplished through the definition of queues that have Forwarding Class (FC), Fair Information Rate (FIR), Committed Information Rate (CIR), Peak Information Rate (PIR), Committed Burst Size (CBS), and Maximum Burst Size (MBS) characteristics.

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 they are assigned to at least one forwarding class and a policy is applied to a service SAP.

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 cannot be modified or deleted. The default SAP ingress policy defines one unicast and one multipoint queue associated with all forwarding classes, with an FIR of zero, a CIR of zero, and a 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, 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 SAP ingress policy. A policy cannot be deleted until it is removed from all services where it is applied.

**Parameters** *policy-id* — The *policy-id* uniquely identifies the policy.

**Values** 1 to 65535

*policy-name* — The *policy-name* uniquely identifies the policy.

**Values** 64 characters maximum

**create** — Required parameter when creating a SAP QoS ingress policy.

**name** *name* — Configures an optional policy name, up to 64 characters in length, which adds a name identifier to a given policy to then use that policy name in configuration references as well as display and use policy names in show commands throughout the system. This helps the service provider/administrator to identify and manage sap-ingress policies within the SR OS platforms.

All sap-ingress policies are required to assign a policy ID to initially create a policy. However, either the policy ID or the policy name can be used to identify and reference a given policy once it is initially created

If a name is not specified at creation time, then SR OS assigns a string version of the *policy-id* as the name.

**Values** 64 characters maximum

## default-fc

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>default-fc</b> <i>fc-name</i><br><b>no default-fc</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>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 subclass 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 enqueueing priority for the traffic can be overridden as well.</p> <p>The default forwarding class is best effort (be). The <b>default-fc</b> settings are displayed in the <b>show configuration</b> and <b>save</b> output regardless of inclusion of the <b>detail</b> keyword.</p>                                                                                            |
| <b>Default</b>     | default-fc "be"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <p><i>fc-name</i> — Specify the forwarding class name for the queue. The value given for <i>fc-name</i> must be one of the predefined forwarding classes in the system.</p> <p>The subclass-name parameter is optional and used with the fc-name parameter to define a preexisting subclass. The fc-name and subclass-name parameters must be separated by a period (dot). If subclass-name does not exist in the context of fc - name, an error will occur. If subclass-name is removed using the <b>no fc fc-name.subclass-name force</b> command, the <b>default-fc</b> command will automatically drop the <i>subclass-name</i> and only use <i>fc-name</i> (the parent forwarding class for the subclass) as the forwarding class.</p> <p><b>Values</b></p> <p>fc: <i>class[.subclass]</i></p> <p>class: be, l2, af, l1, h2, ef, h1, nc</p> <p>subclass: 29 characters max</p> |

## default-priority

|                |                                                                    |
|----------------|--------------------------------------------------------------------|
| <b>Syntax</b>  | <b>default-priority</b> {high   low}<br><b>no default-priority</b> |
| <b>Context</b> | config>qos>sap-ingress                                             |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command configures the default enqueueing priority for all packets received on an ingress SAP using this policy. To change the default priority for the policy, the <b>fc-name</b> must be defined whether it is being changed or not.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>     | default-priority low                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><b>high</b> — Setting the enqueueing parameter to <b>high</b> for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.</p> <p><b>low</b> — Setting the enqueueing parameter to <b>low</b> for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.</p> |

## dot1p

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dot1p</b> <i>dot1p-priority</i> [ <b>fc</b> <i>fc-name</i> ] [ <b>priority</b> { <b>low</b>   <b>high</b> }]<br><b>no dot1p</b> <i>dot1p-priority</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command explicitly sets the forwarding class or subclass or enqueueing priority when a packet is marked with a <i>dot1p-priority</i> specified. Adding a dot1p rule on the policy forces packets that match the <i>dot1p-priority</i> specified to override the forwarding class and enqueueing 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 enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing priority derived from earlier matches in the classification hierarchy.</p> <p>The <i>dot1p-priority</i> 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 QoS behavior.</p> <p>The <b>no</b> 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.</p> |
| <b>Parameters</b>  | <p><i>dot1p-priority</i> — 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 <i>dot1p-value</i>, the previous forwarding class and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.</p> <p>A maximum of eight dot1p rules are allowed on a single policy.</p> <p><b>Values</b>      0 to 7</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**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.

**priority** — This parameter overrides the default enqueueing 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 enqueueing 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 enqueueing priority is inherited based on previous rule matches.

**Default** Inherits the priority defined by the default-priority statement.

**high** — This parameter is used in conjunction with the **priority** parameter. Setting the enqueueing parameter to **high** for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**low** — This parameter is used in conjunction with the **priority** parameter. Setting the enqueueing parameter to **low** for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

## dscp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dscp</b> <i>dscp-name</i> [ <i>dscp-name</i> ] <b>fc</b> <i>fc-name</i> [ <b>priority</b> { <b>low</b>   <b>high</b> }]<br><b>no dscp</b> <i>dscp-name</i> [ <i>dscp-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | This command explicitly sets the forwarding class or subclass or enqueueing priority when a packet is marked with the DiffServ Code Point (DSCP) value contained in the <i>dscp-name</i> . A list of up to eight <i>dscp-names</i> can be entered on a single command. The lists of <i>dscp-names</i> within the configuration are managed by the system to ensure that each list does not exceed eight names. Entering more than eight <i>dscp-names</i> with the same parameters ( <b>fc</b> , <b>priority</b> ) will result in multiple lists being created. Conversely, multiple lists with the same parameters ( <b>fc</b> , <b>priority</b> ) are merged and the lists repacked to a maximum of eight 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. 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 enqueueing 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 enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing 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 this 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 eight 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 enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.

A maximum of 64 DSCP rules are allowed on a single policy and a maximum of eight *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 following command:

```
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 |



## dynamic-policer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dynamic-policer</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>qos>sap-egress<br>config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command enables the context in which common properties for dynamic-policers can be configured. Dynamic policers are instantiated and terminated on demand due to an action request submitted by the policy server (for example, via Gx interface). The actions types behind dynamic policers are typically related to rate-limiting or volume monitoring. The dynamic-policers can be instantiated on demand at any time during the lifetime of the sla-profile instance. |

## cbs

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cbs size [bytes   kilobytes]</b><br><b>no cbs</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>qos>sap-ingress>dyn-policer<br>config>qos>sap-egress>dyn-policer<br>config>qos>sap-ingress>policer<br>config>qos>sap-egress>policer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command configures 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.</p> <p>The policer's <b>cbs</b> size defined in the QoS policy may be overridden on an <b>sla-profile</b> or SAP where the policy is applied.</p> <p>The <b>no</b> form of this command returns the policer to its default CBS size.</p> <p>By default, the CBS is 16 Mbytes when CIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured CBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max CIR capped to 3968 kbytes, with a minimum of 256 bytes.</p> |
| <b>Parameters</b>  | <p><b>size [bytes   kilobytes]</b> — Specifies an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional <b>byte</b> and <b>kilobyte</b> keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.</p> <p><b>Values</b> 0 to 16777216 or default</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

---

mbs

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mbs</b> { <i>size</i> [ <b>bytes</b>   <b>kilobytes</b> ]   <b>default</b> }<br><b>no mbs</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>qos>sap-ingress>dyn-policer<br>config>qos>sap-egress>dyn-policer<br>config>qos>sap-ingress>policer<br>config>qos>sap-egress>policer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command is used to configure the policer's PIR leaky bucket's high-priority violate threshold. The <b>high-prio-only</b> command is applied to the MBS value to derive the bucket's low-priority violate threshold. For ingress, trusted in-profile packets and untrusted high-priority packets use the policer's high-priority violate threshold while trusted out-of-profile and untrusted low-priority packets use the policer's low-priority violate threshold. At egress, inplus-profile, and in-profile packets use the policer's high-priority violate threshold and out-of-profile packets use the policer's low-priority violate threshold. Exceed-profile packets are discarded unless <b>enable-exceed-pir</b> is configured, in which case they are forwarded.</p> <p>The PIR bucket's violate threshold represents 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 <b>high-prio-only</b> is available for the higher priority traffic.</p> <p>The policer's MBS size defined in the QoS policy may be overridden on an SLA profile or SAP where the policy is applied.</p> <p>The <b>no</b> form of this command returns the queue to its default MBS size.</p> <p>By default, the MBS is 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes.</p> |
| <b>Parameters</b>  | <p><b>size</b> [<b>bytes</b>   <b>kilobytes</b>] — The <b>size</b> parameter is required when specifying <b>mbs</b> and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional <b>bytes</b> and <b>kilobytes</b> keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.</p> <p><b>Values</b>      0 – 16777216 or <b>default</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## packet-byte-offset

|               |                                                                                                                               |
|---------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>packet-byte-offset</b> { <b>add</b> <i>add-bytes</i>   <b>subtract</b> <i>sub- bytes</i> }<br><b>no packet-byte-offset</b> |
|---------------|-------------------------------------------------------------------------------------------------------------------------------|

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | <pre>config&gt;qos&gt;sap-ingress&gt;dyn-policer config&gt;qos&gt;sap-egress&gt;dyn-policer config&gt;qos&gt;sap-ingress&gt;policer config&gt;qos&gt;sap-egress&gt;policer</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>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 <b>packet-byte-offset</b> command is meant to be an arbitrary mechanism that 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.</p> <p>When child policers are adding to or subtracting from the size of each packet, the parent policer's <b>min-thresh-separation</b> value should also need to be modified by the same amount.</p> <p>The policer's <b>packet-byte-offset</b> defined in the QoS policy may be overridden on an <b>sla-profile</b> or SAP where the policy is applied. Packet byte offset settings are not included in the applied rate when (queue) frame-based accounting is configured and the policer is managed by HQoS; however, the offsets are applied to the statistics.</p> <p>The <b>no</b> form of this command is used to remove per packet size modifications from the policer.</p> |
| <b>Parameters</b>  | <p><b>add</b> <i>add-bytes</i> — The <b>add</b> keyword is mutually exclusive to the <b>subtract</b> keyword. Either <b>add</b> or <b>subtract</b> must be specified. When <b>add</b> is defined, the corresponding bytes parameter specifies the number of bytes that is added to 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 increased by the amount being added to the size of each packet.</p> <p><b>Values</b>      0 to 31</p> <p><b>subtract</b> <i>sub-bytes</i> — The <b>subtract</b> keyword is mutually exclusive to the <b>add</b> keyword. Either <b>add</b> or <b>subtract</b> must be specified. When <b>subtract</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. The minimum resulting packet size used by the system is 1 byte.</p> <p><b>Values</b>      1 to 64</p>                           |

## parent

|                |                                                                                                                                           |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <pre>parent <i>arbiter-name</i> [<b>weight</b> <i>weight-within-level</i>] [<b>level</b> <i>level</i>] no parent</pre>                    |
| <b>Context</b> | <pre>config&gt;qos&gt;sap-ingress&gt;dyn-policer config&gt;qos&gt;sap-egress&gt;dyn-policer config&gt;qos&gt;sap-ingress&gt;policer</pre> |

```
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 or multiservice site context. To create a policer control hierarchy on an ingress or egress SAP context, a **policer-control-policy** must be applied to the SAP. When 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 is cleared and a resulting trap is 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. This functionality is available only for the 7450 ESS and 7750 SR.

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 or multiservice site 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.

When 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.
- 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**

## range

- Syntax** **range start-entry policer-id count count**  
**no range**
- Context** config>qos>sap-egress>dyn-policer  
config>qos>sap-ingress>dyn-policer
- Description** This command defines the range of ids for dynamic policers that are created via Gx interface.
- The **no** form of this command disables creation of dynamic policers via Gx interface, resulting in a Gx rule instantiation failure, which is the default.
- Default** no range
- Parameters** **start-entry policer-id** — Specifies the lowest entry in the range.
- Values** 1 to 63
- count count** — Specifies the number of entries in the range.
- Values** 1 to 63

## stat-mode

|                |                                                                                                                                                                                                                                                                                                                                                  |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>stat-mode</b> { <b>no-stats</b>   <b>minimal</b>   <b>offered-profile-no-cir</b>   <b>offered-priority-no-cir</b>   <b>offered-profile-cir</b>   <b>offered-priority-cir</b>   <b>offered-total-cir</b>   <b>offered-limited-profile-cir</b>   <b>offered-profile-capped-cir</b>   <b>offered-limited-capped-cir</b> }<br><b>no stat mode</b> |
| <b>Context</b> | config>qos>sap-ingress>dyn-policer<br>config>qos>sap-ingress>policer                                                                                                                                                                                                                                                                             |

**Description** This command is used to configure the forwarding plane counters that allow offered, forwarded, and dropped accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, uncolored, 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, for example, 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 that prevents any packet accounting, the use of the policer's **parent** command requires that the policer's **stat-mode** 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. When 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. The total/allocated/free stats can be viewed by using the **tools dump resource-usage card fp** 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 ingress policer stat-modes are described in [Table 26](#).

**Table 26 Ingress Policer Stat Mode Summary**

| Stat Mode | Stat Resources | Traffic Counters (Packet/Octets) |                       | Comments |
|-----------|----------------|----------------------------------|-----------------------|----------|
|           |                | Offered                          | Dropped/<br>Forwarded |          |
| no-stats  | 0              | None                             | None                  | —        |



**Table 26 Ingress Policer Stat Mode Summary (Continued)**

| Stat Mode                   | Stat Resources | Traffic Counters (Packet/Octets)  |                                                      | Comments                                                                                                                                                                                                             |
|-----------------------------|----------------|-----------------------------------|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                             |                | Offered                           | Dropped/<br>Forwarded                                |                                                                                                                                                                                                                      |
| Minimal                     | 1              | Single counter entering policer   | Single counter for dropped/forwarded exiting policer | —                                                                                                                                                                                                                    |
| offered-profile-no-cir      | 2              | In/out entering policer           | In/out entering policer                              | Intended for when the policer does not change the profile of packets. Includes only in-profile and out-of-profile.                                                                                                   |
| offered-priority-no-cir     | 2              | High/low entering policer         | High/low entering policer                            | Intended for when only packet priority stats are required.                                                                                                                                                           |
| offered-profile-cir         | 4              | In/out/uncolored entering policer | In/out exiting policer                               | Intended for when the policer can change the profile of packets to in-profile and out-of-profile.                                                                                                                    |
| offered-priority-cir        | 4              | High/low entering policer         | In/out exiting policer                               | Intended for when packet priority entering the policer and profile exiting the policer is required.                                                                                                                  |
| offered-total-cir           | 2              | Single counter entering policer   | In/out exiting policer                               | —                                                                                                                                                                                                                    |
| offered-limited-profile-cir | 3              | Out/uncolored entering policer    | In/out exiting policer                               | Intended for when the policer can change the profile of packet to in-profile and out-of-profile. The information is limited compared to offered-profile-capped-cir with the benefit of using one less stat resource. |

**Table 26** Ingress Policer Stat Mode Summary (Continued)

| Stat Mode                  | Stat Resources | Traffic Counters (Packet/Octets)  |                        | Comments                                                                                                                                                                         |
|----------------------------|----------------|-----------------------------------|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                            |                | Offered                           | Dropped/<br>Forwarded  |                                                                                                                                                                                  |
| offered-profile-capped-cir | 5              | In/out/uncolored entering policer | In/out exiting policer | Intended for when the policer has profile-capped configured.                                                                                                                     |
| offered-limited-capped-cir | 4              | In/uncolored entering policer     | In/out exiting policer | Intended for when the policer has profile-capped configured. The information is limited compared to offered-profile-capped-cir with the benefit of using one less stat resource. |

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, dropped, and forwarded 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, no statistics are generated.

**minimal** — Counter resource allocation: 1

This stat-mode provides the minimal accounting resource usage and counter information, and includes the total offered, dropped and forwarded packet and octet counters for traffic entering (offered) and exiting (dropped/forwarded) the policer.

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 in-profile or out-of-profile 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 displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 27](#).

**Table 27 Ingress Accounting Statistics Collected in minimal stat-mode**

| Show Output | Accounting Statistics Collected |                     |
|-------------|---------------------------------|---------------------|
|             | Field                           | Field Description   |
| Off. All    | apo                             | AllPacketsOffered   |
|             | aoo                             | AllOctetsOffered    |
| Dro. All    | apd                             | AllPacketsDropped   |
|             | aod                             | AllOctetsDropped    |
| For. All    | apf                             | AllPacketsForwarded |
|             | aof                             | AllOctetsForwarded  |

**offered-profile-no-cir** — Counter resource allocation: 2

This **stat-mode** provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering the policer.

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 premarked (and trusted) packets. It is expected that, in this instance, a CIR rate will not be defined since all packets are already premarked. This mode does not prevent the policer from receiving untrusted (color undefined) traffic nor does it prevent the policer from being configured with a CIR rate.

This mode is intended to be used without profile-capped configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 28](#).

**Table 28** Ingress Accounting Statistics Collected in offered-profile-no-cir stat-mode

| Show Output  | Accounting Statistics Collected |                              |
|--------------|---------------------------------|------------------------------|
|              | Field                           | Field Description            |
| Off. InProf  | ipo                             | InProfilePacketsOffered      |
|              | ioo                             | InProfileOctetsOffered       |
| Off. OutProf | opo                             | OutOfProfilePacketsOffered   |
|              | ooo                             | OutOfProfileOctetsOffered    |
| Dro. InProf  | ipd                             | InProfilePacketsDropped      |
|              | iod                             | InProfileOctetsDropped       |
| Dro. OutProf | opd                             | OutOfProfilePacketsDropped   |
|              | ood                             | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                             | InProfilePacketsForwarded    |
|              | iof                             | InProfileOctetsForwarded     |
| For. OutProf | opf                             | OutOfProfilePacketsForwarded |
|              | oof                             | OutOfProfileOctetsForwarded  |

**offered-priority-no-cir** — Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the packet priority of traffic entering the policer.

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 untrusted 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 premarked in-profile or out-of-profile nor does it prevent the policer from being configured with a CIR rate.

This mode is intended to be used without profile-capped configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 29](#).

**Table 29 Ingress Accounting Statistics Collected in offered-priority-no-cir stat-mode**

| Show Output  | Accounting Statistics Collected |                              |
|--------------|---------------------------------|------------------------------|
|              | Field                           | Field Description            |
| Off. HiPrio  | hpo                             | HighPriorityPacketsOffered   |
|              | hoo                             | HighPriorityOctetsOffered    |
| Off. LowPrio | lpo                             | LowPriorityPacketsOffered    |
|              | loo                             | LowPriorityOctetsOffered     |
| Dro. HiPrio  | hpd                             | HighPriorityPacketsDropped   |
|              | hod                             | HighPriorityOctetsDropped    |
| Dro. LowPrio | lpd                             | LowPriorityPacketsDropped    |
|              | lod                             | LowPriorityOctetsDropped     |
| For. HiPrio  | hpf                             | HighPriorityPacketsForwarded |
|              | hof                             | HighPriorityOctetsForwarded  |
| For. LowPrio | lpf                             | LowPriorityPacketsForwarded  |
|              | lof                             | LowPriorityOctetsForwarded   |

**offered-profile-cir** — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises hard in/out and uncolored traffic. The offered counters cover traffic explicitly profiled to in-profile, traffic explicitly profiled to out-of-profile, and traffic that has not been explicitly profiled at ingress (uncolored).

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 untrusted 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 untrusted markings. If trusted in-profile packets are not being received, the **offered-limited-profile-cir** stat-mode could be used instead, which has the benefit of using a reduced number of stat resources.

This mode is intended to be used without **profile-capped** configured within the policer as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 30](#).

**Table 30 Ingress Accounting Statistics Collected in offered-profile-cir stat-mode**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. InProf  | ipo                        | InProfilePacketsOffered      |
|              | ioo                        | InProfileOctetsOffered       |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered   |
|              | ooo                        | OutOfProfileOctetsOffered    |
| Off. Uncolor | ucp                        | UncoloredPacketsOffered      |
|              | uco                        | UncoloredOctetsOffered       |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-priority-cir** — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the priority of traffic entering the policer and the profile exiting the policer.

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 untrusted 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.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 31](#).

**Table 31** Ingress Accounting Statistics Collected in offered-priority-cir stat-mode

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. HiPrio  | hpo                        | HighPriorityPacketsOffered   |
|              | hoo                        | HighPriorityOctetsOffered    |
| Off. LowPrio | lpo                        | LowPriorityPacketsOffered    |
|              | loo                        | LowPriorityOctetsOffered     |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-total-cir** — Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter.

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 untrusted 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 untrusted packets.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 32](#).

**Table 32**      **Ingress Accounting Statistics Collected in offered-total-cir stat-mode**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. All     | apo                        | AllPacketsOffered            |
|              | aoo                        | AllOctetsOffered             |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-limited-profile-cir** — Counter resource allocation: 3

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard out and uncolored. The offered counters cover traffic explicitly profiled to out-of-profile and traffic that has not been explicitly profiled at ingress (Uncolor). The traffic explicitly profiled to in-profile is counted with the uncolored traffic.

The **offered-limited-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 untrusted packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile packets. If trusted in-profile packets are not being received, the **offered-limited-profile-cir** is preferred over **offered-profile-cir** because it uses a reduced number of stat resources.



This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 33](#).

**Table 33** Ingress Accounting Statistics Collected in offered-limited-profile-cir stat-mode

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered   |
|              | ooo                        | OutOfProfileOctetsOffered    |
| Off. Uncolor | ucp                        | UncoloredPacketsOffered      |
|              | uco                        | UncoloredOctetsOffered       |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-profile-cir** — Counter resource allocation: 4

**offered-profile-capped-cir** — Counter resource allocation: 5

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard in/out and uncolored. The offered counters cover traffic explicitly profiled to in-profile, traffic explicitly profiled to out-of-profile, and traffic that has not been explicitly profiled at ingress (Uncolor).

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 is that one of the counting resources in the forwarding plane and traffic manager will not be used.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 34](#).

**Table 34** Ingress Accounting Statistics Collected in offered-profile-capped-cir stat-mode

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. InProf  | ipo                        | InProfilePacketsOffered      |
|              | ioo                        | InProfileOctetsOffered       |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered   |
|              | ooo                        | OutOfProfileOctetsOffered    |
| Off. Uncolor | ucp                        | UncoloredPacketsOffered      |
|              | uco                        | UncoloredOctetsOffered       |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-limited-capped-cir** — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed resulting in the traffic entering the policer comprising of hard in/out and uncolored. The offered counters cover in-profile traffic and traffic that has not been explicitly profiled at ingress (Uncolor). The traffic explicitly profiled to out-of-profile is counted with the uncolored traffic.

When **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 **soft in-profile** with **profile in** (InProf) and **profile out** (OutProf) with **soft-out-of-profile** (Uncolor) and eliminates the “offered undefined” statistic. If trusted out-of-profile packets are not being received, the **offered-limited-capped-cir** is preferred over **offered-profile-capped-cir** because it uses a reduced number of stat resources.

This mode is intended to be used with **profile-capped** configured within the policer.

The impact of using **offered-limited-capped-cir** stat-mode while **profile-capped** mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 35](#).

**Table 35 Ingress Accounting Statistics Collected in offered-limited-capped-cir stat-mode**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. InProf  | ipo                        | InProfilePacketsOffered      |
|              | ioo                        | InProfileOctetsOffered       |
| Off. Uncolor | ucp                        | UncoloredPacketsOffered      |
|              | uco                        | UncoloredOctetsOffered       |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

---

fc

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fc</b> <i>fc-name</i> [create]<br><b>no</b> <b>fc</b> <i>fc-name</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>The <b>fc</b> command creates a class or subclass instance of the forwarding class <i>fc-name</i>. When the <i>fc-name</i> is created, classification actions can be applied and the subclass can be used in match classification criteria. Attempting to use an undefined subclass in a classification command will result in an execution error and the command will fail.</p> <p>The <b>no</b> form of this command removes all the explicit queue mappings for <i>fc-name</i> forwarding types. The queue mappings revert to the default queues for <i>fc-name</i>. To successfully remove a subclass, all associations with the subclass in the classification commands within the policy must first be removed or diverted to another forwarding class or subclass.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <p><i>fc-name</i> — The parameter subclass-name is optional and must be defined using a dot separated notation with a preceding valid system-wide forwarding class name. Creating a subclass 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 subclass. A subclass must always be specified with its preceding forwarding class name. When a forwarding class is created or specified without the optional subclass, the parent forwarding class is assumed.</p> <p>Within the SAP ingress QoS policy, up to 56 subclasses may be created. Each of the 56 subclasses may be created within any of the eight parental forwarding classes. When the limit of 56 is reached, any further subclass creations will fail and the subclass will not exist.</p> <p>Successfully creating a subclass places the CLI within the context of the subclass for further subclass parameter definitions. Within the subclass context, commands may be executed that define subclass priority (within the parent forwarding class queue mapping), subclass color aware profile settings, subclass in-profile and out-of-profile precedence or DSCP markings.</p> <p>The <i>subclass-name</i> parameter is optional and used with the <i>fc-name</i> parameter to define a pre-existing subclass. The <i>fc-name</i> and <i>subclass-name</i> parameters must be separated by a period (.). If <i>subclass-name</i> does not exist in the context of <i>fc-name</i>, an error will occur. If <i>subclass-name</i> is removed using the <b>no fc fc-name.subclass-name force</b> command, the <b>default-fc</b> command will automatically drop the <i>subclass-name</i> and only use <i>fc-name</i> (the parent forwarding class for the subclass) as the forwarding class.</p> <p><b>Values</b></p> <p>fc:            <i>class[.subclass]</i></p> <p>                    <i>class</i>: be, l2, af, l1, h2, ef, h1, nc</p> <p>                    <i>subclass</i>: 29 characters max</p> |

**create** — Required parameter when creating a SAP QoS ingress policy forwarding class.

## broadcast-policer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>broadcast-policer</b> <i>policer-id</i> [ <b>fp-redirect-group</b> ]<br><b>no broadcast-policer</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>qos>sap-ingress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | Within a <b>sap-ingress</b> QoS policy forwarding class context, the <b>broadcast-policer</b> 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 <b>sap-ingress</b> 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 multiservice site, or ingress policing is not supported on the port associated with the SAP or subscriber or multiservice site, 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 multiservice site 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 subclasses 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 or multiservice site 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 to 63
- fp-redirect-group** — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

## broadcast-queue

- Syntax** **broadcast-queue** *queue-id* [**group** *queue-group-name*]  
**no broadcast queue**
- Context** config>qos>sap-ingress>fc
- 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. When 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 this 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 that 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.

All received DE = 0 frames that are classified into this parent forwarding class or any of its subclasses 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 that 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 their priority automatically set to low while DE = 0 frames have their priority set to high. Furthermore, DE = 1 frames have the 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 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 that 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 that is part of an Fpipe service. It can also be used on the ingress of an Epipe or VPLS SAP.

The **no** form of this command disables the color profile mode of operation on all SAPs to which 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 not the QoS Policy Propagation via BGP.

The forwarding class and/or forwarding subclass can be overridden.

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.

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

## in-remark

**Syntax** **in-remark** {**dscp** *dscp-name* | **prec** *ip-prec-value*}  
**no in-remark**

**Context** config>qos>sap-ingress>fc

**Description** This command is used in a SAP ingress QoS policy to define an explicit in-profile remark action for a forwarding class or subclass. 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 or subscriber interface (in an IES or VPRN). 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 subclass. It is possible for a packet to match a classification command that maps the packet to a particular forwarding class or subclass, only to have a more explicit (higher priority match) override the association. Only the highest priority match forwarding class or subclass 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. [Table 36](#) 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.

**Table 36 Effect of In-Remark Command on Received SAP Ingress Packets**

| 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)              |



**Table 36 Effect of In-Remark Command on Received SAP Ingress Packets (Continued)**

| SAP Ingress Packet State           | in-remark Command Effect                       |
|------------------------------------|------------------------------------------------|
| 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 this command disables ingress remarking of in-profile packets classified to the forwarding class or subclass.

- Parameters** **dscp** *dscp-name* — Specifies that the matching packet's DSCP bits should be overridden with the value represented by *dscp-name*.  
The *dscp-name* parameter is a 6-bit value. 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
- prec** *ip-prec-value* — Specifies that the matching packet's precedence bits should be overridden with the value represented by *ip-prec-value*.
- Values** 0 to 7

## 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 multiservice site, or ingress policing is not supported on the port associated with the SAP or subscriber or multiservice site, the initial forwarding class forwarding type mapping will fail.

The **multicast-policer** command is ignored for instances of the policer applied to SAPs subscribers or multiservice site 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 subclasses 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 or multiservice site 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 to 63

**fp-redirect-group** — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

## multicast-queue

**Syntax**           **multicast-queue queue-id [group queue-group-name]**  
**no multicast-queue**

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>qos>sap-ingress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command overrides the default multicast forwarding type queue mapping for <b>fc</b> <i>fc-name</i>. The specified <i>queue-id</i> must exist within the policy as a multipoint queue before the mapping can be made. When the forwarding class mapping is executed, all multicast traffic on a SAP using this policy is forwarded using the <i>queue-id</i>.</p> <p>The multicast forwarding type includes the <b>unknown</b> unicast forwarding type and the <b>broadcast</b> 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.</p> <p>The <b>no</b> form of this command sets the multicast forwarding type <i>queue-id</i> back to the default queue for the forwarding class. If the <b>broadcast</b> and <b>unknown</b> forwarding types were not explicitly defined to a multipoint queue, they will also be set back to the default multipoint queue (queue 11).</p> |
| <b>Parameters</b>  | <p><i>queue-id</i> — The <i>queue-id</i> parameter specified must be an existing, multipoint queue defined in the config&gt;qos&gt;sap-ingress context.</p> <p><b>Values</b> Any valid multipoint queue-ID in the policy including 2 through 32.</p> <p><b>Default</b> 11</p> <p><b>group</b> <i>queue-group-name</i> — 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 <i>queue-group-name</i> are configured in the <b>config&gt;qos&gt;queue-group-templates</b> egress and ingress contexts.</p>                                                                                                                                                                                                                                                                                        |

## out-remark

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>out-remark {dscp dscp-name   prec ip-prec-value}</b><br><b>no out-remark</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-ingress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command is used in a SAP ingress QoS policy to define an explicit out-of-profile remark action for a forwarding class or subclass. 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 or subscriber interface (in an IES or VPRN). When the policy is applied to a Layer 2 SAP (for example, Epipe or VPLS), the remarking definitions are silently ignored.</p> <p>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 subclass. It is possible for a packet to match a classification command that maps the packet to a particular forwarding class or subclass, only to have a more explicit (higher priority match) override the association. Only the highest priority match forwarding class or subclass association will drive the out-of-profile marking.</p> |

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. [Table 37](#) describes 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 37 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.

An explicit dscp name or precedence value must be specified for out-of-profile remarking to be applied.

The **no** form of this command disables ingress remarking of out-of-profile packets classified to the forwarding class or subclass.

**Default** no out-remark

**Parameters** **dscp** *dscp-name* — Specifies that the matching packet's DSCP bits should be overridden with the value represented by *dscp-name*.  
The *dscp-name* parameter is a 6-bit value. 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

**prec** *ip-prec-value* — Specifies that the matching packet's precedence bits should be overridden with the value represented by *ip-prec-value*.

The value specified by *ip-prec-value* is used to overwrite the precedence bits within a matching routed packets IP header ToS field.

**Values** 0 to 7

## policer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>policer</b> <i>policer-id</i> [ <b>fp-redirect-group</b> ]<br><b>no policer</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>qos>sap-ingress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>Within a sap-ingress QoS policy forwarding class context, the <b>policer</b> 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 <b>sap-ingress</b> 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 <b>queue</b> <i>queue-id</i> or <b>queue</b> <i>queue-id</i> <b>group</b> <i>ingress-queue-group</i> commands) or an ingress policer (<b>policer</b> <i>policer-id</i>). The <b>queue</b> and <b>policer</b> 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 <b>policer</b> <i>policer-id</i> 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.</p> |

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 multiservice site or ingress policing is not supported on the port associated with the SAP or subscriber or multiservice site, 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 subclasses 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 subscriber or multiservice sites 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 to 63

**fp-redirect-group** — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

## profile

**Syntax** **profile {in | out}**  
**no profile**

**Context** config>qos>sap-ingress>fc

**Description** This command places a forwarding class or subclass 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 subclass) 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 this command removes an explicit in-profile or out-of-profile configuration on a forwarding class or subclass.

**Default** **no profile** — The default profile state of a forwarding class or subclass 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 subclass.

**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.

## queue

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>queue</b> <i>queue-id</i> [ <b>group</b> <i>queue-group-name</i> ]<br><b>no queue</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>qos>sap-ingress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command overrides the default queue mapping for <b>fc</b> <i>fc-name</i>. The specified <i>queue-id</i> must exist within the policy before the mapping can be made. When the forwarding class mapping is executed, all traffic classified to <i>fc-name</i> on a SAP using this policy.</p> <p>The <b>no</b> form of this command sets the <i>queue-id</i> back to the default queue for the forwarding class.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>     | queue 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <p><i>queue-id</i> — Specifies the SAP egress <i>queue-id</i> to be associated with the forwarding class. The <i>queue-id</i> must be an existing queue defined in sap-egress policy-id.</p> <p><b>Values</b>      1 — 8</p> <p><b>Default</b>      1</p> <p><b>group</b> <i>queue-group-name</i> — This optional parameter is used to redirect the forwarding type within the forwarding class to the specified <i>queue-id</i> within the <i>queue-group-name</i>. 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 <i>queue-group-name</i> are configured in the <b>config&gt;qos&gt;queue-group-templates</b> 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.</p> |

## unknown-policer

|                |                                                                                                    |
|----------------|----------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>unknown-policer</b> <i>policer-id</i> [ <b>fp-redirect-group</b> ]<br><b>no unknown-policer</b> |
| <b>Context</b> | 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 multiservice site, or ingress policing is not supported on the port associated with the SAP or subscriber or multiservice site, 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 multiservice site 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 subclasses 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 or multiservice site 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 to 63



**fp-redirect-group** — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

## unknown-queue

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>unknown-queue</b> <i>queue-id</i> [ <b>group</b> <i>queue-group-name</i> ]<br><b>no unknown-queue</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>qos>sap-ingress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command overrides the default unknown unicast forwarding type queue mapping for <b>fc</b> <i>fc-name</i>. The specified <i>queue-id</i> must exist within the policy as a multipoint queue before the mapping can be made. When the forwarding class mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the <i>queue-id</i>.</p> <p>The unknown forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.</p> <p>The <b>no</b> form of this command sets the unknown forwarding type <i>queue-id</i> back to the default of tracking the multicast forwarding type queue mapping.</p>                                                 |
| <b>Parameters</b>  | <p><i>queue-id</i> — Specifies an existing multipoint queue defined in the <b>config&gt;qos&gt;sap-ingress</b> context.</p> <p><b>Values</b> Any valid multipoint <i>queue-id</i> in the policy including 2 through 32.</p> <p><b>Default</b> 11</p> <p><b>group</b> <i>queue-group-name</i> — 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 <i>queue-group-name</i> are configured in the <b>config&gt;qos&gt;queue-group-templates</b> egress and ingress contexts.</p> |

## ip-criteria

|                    |                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ip-criteria</b>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>qos>sap-egress<br>config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>IP criteria-based SAP ingress or egress policies are used to select the appropriate ingress or egress queue or policer and corresponding forwarding class and packet profile for matched traffic.</p> <p>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.</p> |

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. When IP criteria entries are removed from a SAP ingress or egress policy, the IP criteria is removed from all services where that policy is applied.

## entry

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>entry</b> <i>entry-id</i> [ <b>create</b> ]<br><b>no entry</b> <i>entry-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>qos>sap-egress>ip-criteria<br>config>qos>sap-ingress>ip-criteria<br>config>qos>sap-egress>ipv6-criteria<br>config>qos>sap-ingress>ipv6-criteria<br>config>qos>sap-ingress>mac-criteria                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command is used to create or edit an IP, IPv6, or MAC criteria entry for the policy. Multiple entries can be created using unique <i>entry-id</i> numbers.</p> <p>The list of flow criteria is evaluated in a top-down manner 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.</p> <p>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.</p> <p>The <b>no</b> 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.</p> |
| <b>Parameters</b>  | <p><i>entry-id</i> — The <i>entry-id</i>, expressed as an integer, uniquely identifies a match criterion and the corresponding action. It is recommended that multiple entries be given <i>entry-ids</i> in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.</p> <p>An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword <b>action fc</b> <i>fc-name</i> for it to be considered complete. Entries without the action keyword will be considered incomplete and, therefore, will be rendered inactive.</p> <p><b>Values</b>      1 to 65535</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**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.

## action

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>action</b> [ <b>fc</b> <i>fc-name</i> ] [ <b>priority</b> { <b>high</b>   <b>low</b> }] [ <b>policer</b> <i>policer-id</i> ]<br><b>no action</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>qos>sap-ingress>ip-criteria>entry<br>config>qos>sap-ingress>ipv6-criteria>entry<br>config>qos>sap-ingress>mac-criteria>entry                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This mandatory command associates the forwarding class or enqueueing priority with specific IP, IPv6, or MAC criteria entry ID. The action command supports setting the forwarding class parameter to a subclass. Packets that meet all match criteria within the entry have their forwarding class and enqueueing priority overridden based on the parameters included in the <b>action</b> parameters. When the forwarding class is not specified in the <b>action</b> command syntax, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueueing priority is not specified in the action, a matching packet preserves (or inherits) the existing enqueueing 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 enqueueing priority derived from previous classification rules.

Each time action is executed on a specific entry ID, the previously entered values for **fc** *fc-name* and **priority** are overridden with the newly defined parameters, or inherit previous matches when a parameter is omitted.

The **no** form of this 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.

If no action is specified, the action specified by the **default-fc** command will be used.

**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 ***subclass-name*** parameter is optional and used with the ***fc-name*** parameter to define a pre-existing subclass. The ***fc-name*** and ***subclass-name*** parameters must be separated by a period (.). If ***subclass-name*** does not exist in the context of ***fc-name***, an error will occur. If ***subclass-name*** is removed using the **no fc *fc-name*.*subclass-name* force** command, the **default-fc** command will automatically drop the ***subclass-name*** and only use ***fc-name*** (the parent forwarding class for the subclass) as the forwarding class.

**Values**

```
fc: class[.subclass]
 class: be, l2, af, l1, h2, ef, h1, nc
 subclass: 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 enqueueing 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 enqueueing 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 enqueueing priority is inherited based on previous rule matches.

**Default**    Inherit (When the **priority (high or low)** is not defined, the rule preserves the previous enqueueing priority of the packet)

**high** — The **high** parameter is used in conjunction with the **priority** parameter. Setting the **priority** enqueueing parameter to **high** for a packet increases the likelihood to enqueue the packet when the queue is congested. The enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the queue, the significance of the enqueueing priority is lost.

**low** — The **low** parameter is used in conjunction with the **priority** parameter. Setting the **priority** enqueueing parameter to **low** for a packet decreases the likelihood to enqueue the packet when the queue is congested. The enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing 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 to 63

## match

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>match</b> [ <b>protocol</b> <i>protocol-id</i> ]<br><b>no match</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>qos>sap-egress>ip-criteria>entry<br>config>qos>sap-ingress>ip-criteria>entry                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>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.</p> <p>If more than one match criteria (within one match statement) are configured, all criteria must be satisfied (AND function) before the action associated with the match is executed.</p> <p>A <b>match</b> context can consist of multiple match criteria, but multiple <b>match</b> statements cannot be entered per entry.</p> <p>It is possible that a SAP policy includes the <b>dscp</b> map command, the <b>dot1p</b> 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:</p> <ol style="list-style-type: none"> <li>1. 802.1p bits</li> <li>2. DSCP</li> <li>3. IP quintuple or MAC headers</li> </ol> <p>The <b>no</b> form of this command removes the match criteria for the <i>entry-id</i>.</p>                                                                                                 |
| <b>Parameters</b>  | <p><b>protocol</b> <i>protocol-id</i> — Specifies an IP protocol to be used as a SAP QoS policy match criterion.</p> <p>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)</p> <p><a href="#">Table 38</a> lists the IP protocols and their respective IDs and descriptions.</p> <p><b>Values</b>      The following values apply to the 7750 SR and 7950 XRS:</p> <p>protocol-id: 0 to 255 protocol numbers accepted in decimal, hexadecimal, or binary</p> <p>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</p> <p><b>Values</b>      The following values apply to the 7450 ESS:</p> <p>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</p> <p>* — udp/tcp wildcard</p> |

**Table 38** 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          | OSPF/IGP                                                    |
| 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 38 IP Protocol Names (Continued)**

| Protocol | Protocol ID | Description                       |
|----------|-------------|-----------------------------------|
| stp      | 118         | Schedule Transfer Protocol        |
| ptp      | 123         | Performance Transparency Protocol |
| isis     | 124         | ISIS over IPv4                    |
| crtip    | 126         | Combat Radio Transport Protocol   |
| crudp    | 127         | Combat Radio User Datagram        |

## dscp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dscp</b> <i>dscp-name</i><br><b>no dscp</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>qos>sap-ingress>ip-criteria>entry>match<br>config>qos>sap-egress>ip-criteria>entry>match<br>config>qos>sap-ingress>ipv6-criteria>entry>match<br>config>qos>sap-egress>ipv6-criteria>entry>match                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command configures a DSCP code point to be used as a SAP QoS policy match criterion.<br><br>The <b>no</b> form of this command removes the DSCP match criterion.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>     | no dscp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <i>dscp-name</i> — Specifies a dscp name that has been previously mapped to a value using the <b>dscp-name</b> command. The DiffServ code point can only be specified by its name.<br><br><b>Values</b> 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 |

## dst-ip

|                |                                                                                                                                                               |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>dst-ip</b> { <i>ip-address/mask</i>   <i>ip-address</i> [ <i>ipv4-address-mask</i> ]   <b>ip-prefix-list</b> <i>prefix-list-name</i> }<br><b>no dst-ip</b> |
| <b>Context</b> | config>qos>sap-ingress>ip-criteria>entry>match<br>config>qos>sap-egress>ip-criteria>entry>match                                                               |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command configures a destination address range to be used as a SAP QoS policy match criterion.</p> <p>To match on the IPv4 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 for IPv4.</p> <p>The <b>no</b> form of this command removes the destination IPv4 address match criterion.</p>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>     | no dst-ip                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <p><i>ip-address</i> — Specifies the destination IPv4 address specified in dotted decimal notation.</p> <p><b>Values</b> ip-address: a.b.c.d</p> <p><i>mask</i> — Specify the length in bits of the subnet mask.</p> <p><b>Values</b> 1 to 32</p> <p><i>ipv4-address-mask</i> — Specify the subnet mask in dotted decimal notation.</p> <p><b>Values</b> a.b.c.d (dotted quad equivalent of mask length)</p> <p><b>ip-prefix-list</b> — creates a list of IPv4 prefixes for match criteria in QoS policies. An ip-prefix-list must contain only IPv4 address prefixes.</p> <p><i>prefix-list-name</i> — A string of up to 32 printable ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> |

## dst-port

|                    |                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>dst-port</b> {<b>lt</b>   <b>gt</b>   <b>eq</b>} <i>dst-port-number</i></p> <p><b>dst-port range</b> <i>start end</i></p> <p><b>no dst-port</b></p>                                                                                                                                                |
| <b>Context</b>     | <p>config&gt;qos&gt;sap-ingress&gt;ip-criteria&gt;entry&gt;match</p> <p>config&gt;qos&gt;sap-egress&gt;ip-criteria&gt;entry&gt;match</p> <p>config&gt;qos&gt;sap-ingress&gt;ipv6-criteria&gt;entry&gt;match</p> <p>config&gt;qos&gt;sap-egress&gt;ipv6-criteria&gt;entry&gt;match</p>                    |
| <b>Description</b> | <p>This command configures a destination TCP or UDP port number or port range for a SAP QoS policy match criterion.</p> <p>The <b>no</b> form of this command removes the destination port match criterion.</p>                                                                                          |
| <b>Default</b>     | no dst-port                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <p>{<b>lt</b>   <b>gt</b>   <b>eq</b>} <i>dst-port-number</i> — The TCP or UDP port numbers to match, specified as less than (<b>lt</b>), greater than (<b>gt</b>), or equal to (<b>eq</b>) to the destination port value, specified as a decimal integer.</p> <p><b>Values</b> 1 to 65535 (decimal)</p> |



**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 to 65535 (decimal)

## fragment

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fragment</b> { <b>true</b>   <b>false</b> }<br><b>no fragment</b>                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>qos>sap-ingress>ip-criteria>entry>match<br>config>qos>sap-egress>ip-criteria>entry>match                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command configures fragmented or non-fragmented IP packets as a SAP QoS policy match criterion.<br><br>The <b>no</b> form of this command removes the match criterion and matches all packets regardless of whether they are fragmented or not.                                                                                                                                                                                 |
| <b>Default</b>     | no fragment                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <b>true</b> — 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.<br><br><b>false</b> — 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

|                    |                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>src-ip</b> { <i>ip-address/mask</i>   <i>ip-address [ipv4-address-mask]</i>   <b>ip-prefix-list</b> <i>prefix-list-name</i> }<br><b>no src-ip</b>                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>qos>sap-ingress>ip-criteria>entry>match<br>config>qos>sap-egress>ip-criteria>entry>match                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command configures a source IPv4 address range to be used as an SAP QoS policy match criterion.<br><br>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.<br><br>The <b>no</b> form of this command removes the source IPv4 or IPv6 address match criterion. |
| <b>Default</b>     | no src-ip                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the source IPv4 address specified in dotted decimal notation.<br><br><b>Values</b> ip-address: a.b.c.d                                                                                                                                                                                                                                                                    |

*mask* — Specifies the length in bits of the subnet mask.

**Values** 1 to 32

*ipv4-address-mask* — Specifies the subnet mask in dotted decimal notation.

**Values** a.b.c.d (dotted quad equivalent of mask length)

**ip-prefix-list** — creates a list of IPv4 prefixes for match criteria in QoS policies. An ip-prefix-list must contain only IPv4 address prefixes.

*prefix-list-name* — A string of up to 32 printable ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## src-port

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>src-port</b> { <b>lt</b>   <b>gt</b>   <b>eq</b> } <i>src-port-number</i><br><b>src-port range</b> <i>start end</i><br><b>no src-port</b>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>qos>sap-ingress>ip-criteria>entry>match<br>config>qos>sap-egress>ip-criteria>entry>match<br>config>qos>sap-ingress>ipv6-criteria>entry>match<br>config>qos>sap-egress>ipv6-criteria>entry>match                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures a source TCP or UDP port number or port range for a SAP QoS policy match criterion.</p> <p>The <b>no</b> form of this command removes the source port match criterion.</p>                                                                                                                                                                                                                                                                                            |
| <b>Default</b>     | no src-port                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <p><b>{lt   gt   eq}</b> <i>src-port-number</i> — The TCP or UDP port numbers to match, specified as less than (<b>lt</b>), greater than (<b>gt</b>), or equal to (<b>eq</b>) to the source port value, specified as a decimal integer.</p> <p><b>Values</b> 1 to 65535 (decimal)</p> <p><b>range start end</b> — The range of TCP or UDP port values to match, specified as between the <i>start</i> and <i>end</i> source port values inclusive.</p> <p><b>Values</b> 1 to 65535 (decimal)</p> |

## vxlan-vni

|                |                                                                                                           |
|----------------|-----------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>vxlan-vni</b> [ <b>eq</b> <i>vxlan-vni-id</i>   <b>range</b> <i>start end</i> ]<br><b>no vxlan-vni</b> |
| <b>Context</b> | config>qos>sap-ingress>ip-criteria>entry>match<br>config>qos>sap-ingress>ipv6-criteria>entry>match        |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command configures a VXLAN or VXLAN GPE VNI to be used as a SAP QoS policy match criterion. A range of VNIs to be matched can be specified by including the keyword <b>range</b> with a <i>start</i> and <i>end</i> VNI. This command requires the <b>type</b> to be set to <b>vlan-vni</b> in the related <b>ip-criteria</b> or <b>ipv6-criteria</b> context.</p> <p>See Virtual Network Identifier (VNI) Classification for the list of restrictions for this command.</p>                                                                                                                                                                                                                                                        |
| <b>Default</b>     | no vxlan-vni                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><b>eq vxlan-vni-id</b> — Specifies the VXLAN or VXLAN GPE VNI to be matched in the SAP ingress QoS classification. The VNI can be specified in any of the available formats but is always shown in decimal.</p> <p><b>Values</b> 1 to 16777215 (Decimal)<br/>0x1 to 0xFFFFFFFF (Hexadecimal)<br/>[0b1 to 0b111111111111111111111111] (Binary)</p> <p><b>range start end</b> — Identifies a range of VNIs to be used as matching criteria. The <i>start</i> value must be lower than the <i>end</i> value. The VNI can be specified in any of the available formats but is always shown in decimal.</p> <p><b>Values</b> 1 to 16777215 (Decimal)<br/>0x1 to 0xFFFFFFFF (Hexadecimal)<br/>[0b1 to 0b111111111111111111111111] (Binary)</p> |

## renum

|                    |                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>renum</b> <i>old-entry-id new-entry-id</i>                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-ingress>ip-criteria<br>config>qos>sap-egress>ip-criteria<br>config>qos>sap-ingress>ipv6-criteria<br>config>qos>sap-egress>ipv6-criteria<br>config>qos>sap-ingress>mac-criteria                                                                                                                                                                      |
| <b>Description</b> | <p>This command renumbers existing QoS policy criteria entries to properly sequence policy entries.</p> <p>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.</p> |
| <b>Parameters</b>  | <p><i>old-entry-id</i> — Enter the entry number of an existing entry.</p> <p><b>Values</b> 1 to 65535</p> <p><i>new-entry-id</i> — Enter the new entry number to be assigned to the old entry.</p> <p><b>Values</b> 1 to 65535</p>                                                                                                                                 |

## type

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>type</b> <i>filter-type</i><br><b>no type</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>qos>sap-ingress>ip-criteria<br>config>qos>sap-ingress>ipv6-criteria                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | This command sets the <b>ip-criteria</b> and <b>ipv6-criteria</b> type to control the type of match entries configurable in this context.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>     | type normal                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <i>filter-type</i> — Specifies which type of entries that the <b>ip-criteria</b> and <b>ipv6-criteria</b> statements can contain.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Values</b>      | <p><b>normal</b> — Regular match criteria are allowed; VXLAN VNI matching is not allowed.</p> <p><b>vxlan-vni</b> — Matching is allowed on a VXLAN VNI for VXLAN and VXLAN GPE frames.</p> <p>The type cannot be changed when <b>ip-criteria</b> or <b>ipv6-criteria</b> entries are configured. If there are no <b>ip-criteria</b> or <b>ipv6-criteria</b> entries configured, the type can be changed from <b>vxlan-vni</b> to <b>normal</b>. The type can only be changed from <b>normal</b> to <b>vxlan-vni</b> if there are no <b>ip-criteria</b> or <b>ipv6-criteria</b> entries configured and if the SAP ingress QoS policy has not been applied to any object</p> <p>A SAP ingress QoS policy configured with <b>type vxlan-vni</b> can be applied to any Ethernet SAPs, except for a PW-SAP, B-VPLS SAP, or CCAG SAP, in any applicable service.</p> |

## ipv6-criteria

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ipv6-criteria</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>qos>sap-egress<br>config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>IPv6 criteria-based SAP egress or ingress policies are used to select the appropriate ingress or egress queue or policer and corresponding forwarding class and packet profile for matched traffic.</p> <p>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.</p> <p>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.</p> <p>This requires mixed-mode for the 7450 ESS.</p> |

The **no** form of this command deletes all the entries specified under this node. When ipv6-criteria entries are removed from a SAP ingress policy, the ipv6-criteria is removed from all services where that policy is applied.

## match

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>match</b> [ <i>next-header next-header</i> ]<br><b>no match</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-ingress>ipv6-criteria>entry<br>config>qos>sap-egress>ipv6-criteria>entry                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>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.</p> <p>If more than one match criteria (within one match statement) are configured, all criteria must be satisfied (AND function) before the action associated with the match is executed.</p> <p>A <b>match</b> context can consist of multiple match criteria, but multiple <b>match</b> statements cannot be entered per entry.</p> <p>It is possible that a SAP ingress policy includes the <b>dscp</b> map command, the <b>dot1p</b> 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:</p> <ol style="list-style-type: none"> <li>1. 802.1p bits</li> <li>2. DSCP</li> <li>3. IP quintuple or MAC headers</li> </ol> <p>The <b>no</b> form of this command removes the match criteria for the <i>entry-id</i>.</p> |
| <b>Parameters</b>  | <p><b>next-header</b> <i>next-header</i> — Specifies the next header to match.</p> <p>On the 7750 SR and 7950 XRS, 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).</p> <p><b>Values</b> protocol numbers accepted in decimal, hexadecimal, or binary: 0 to 42, 45 to 49, 52 to 59, 61 to 255</p> <p><b>keywords:</b> 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</p> <p>* — udp/tcp wildcard</p>                                                                                                                                                                                                                                                                                                                                                                                                                               |

## dst-ip

**Syntax** **dst-ip** {*ipv6-address/prefix-length* | *ipv6-address ipv6-address-mask*}

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <b>no dst-ip</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>qos>sap-ingress>ipv6-criteria>entry>match<br>config>qos>sap-egress>ipv6-criteria>entry>match                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command configures a destination address range to be used as a SAP QoS policy match criterion.</p> <p>To match on the IPv6 destination address, specify the address and its associated mask, e.g. 2001:db8:1000::/64.</p> <p>The <b>no</b> form of this command removes the destination IPv6 address match criterion.</p>                                                                                                                                                                                                                                                                               |
| <b>Default</b>     | no dst-ip                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <p><i>ipv6-address</i> — The IPv6 prefix for the IP match criterion in hex digits (applies to the 7750 SR and 7950 XRS).</p> <p><b>Values</b>      ipv6-address x:x:x:x:x:x:x (eight 16-bit pieces)<br/> x:x:x:x:x::d.d.d.d<br/> x: [0 to FFFF]H<br/> d: [0 to 255]D</p> <p><i>prefix-length</i> — The IPv6 prefix length for the ipv6-address expressed as a decimal integer (applies to the 7750 SR and 7950 XRS).</p> <p><b>Values</b>      1 to 128</p> <p><i>mask</i> — Eight 16-bit hexadecimal pieces representing bit match criteria.</p> <p><b>Values</b>      x:x:x:x:x:x:x (eight 16-bit pieces)</p> |

## fragment

|                    |                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fragment {true   false   first-only   non-first-only}</b><br><b>no fragment</b>                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>qos>sap-ingress>ipv6-criteria>entry>match                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command configures fragmented or non-fragmented IPv6 packets as a SAP ingress QoS policy match criterion.</p> <p>The <b>no</b> form of this command removes the match criterion and matches all packets regardless of whether they are fragmented or not.</p>                                                                          |
| <b>Default</b>     | no fragment                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <p><b>true</b> — Specifies to match on all fragmented IPv6 packets. A match will occur for all packets that contain an IPv6 Fragmentation Extension Header.</p> <p><b>false</b> — Specifies to match on all non-fragmented IP packets. Non-fragmented IPv6 packets are packets that do not contain an IPv6 Fragmentation Extension Header.</p> |

**first-only** — Matches if a packet is an initial fragment of the fragmented IPv6 packet.

**non-first-only** — Matches if a packet is a non-initial fragment of the fragmented IPv6 packet.

## src-ip

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>src-ip</b> { <i>ipv6-address/prefix-length</i>   <b>ipv6-address</b> <i>ipv6-address-mask</i> }<br><b>no src-ip</b>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>sap-ingress>ipv6-criteria>entry>match<br>config>qos>sap-egress>ipv6-criteria>entry>match                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command configures a source IPv6 address range to be used as an SAP QoS policy match criterion.<br><br>To match on the source IPv6 address, specify the address and its associated mask, for example, 2001:db8:1000::/64.<br><br>The <b>no</b> form of this command removes the source IPv6 address match criterion.                                                                                                                                                                                                                           |
| <b>Default</b>     | no src-ip                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <i>ipv6-address</i> — Specifies the IPv6 prefix for the IP match criterion in hex digits.<br><br><b>Values</b> <i>ipv6-address</i> : x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:x.d.d.d.d<br>x: [0 to FFFF]H<br>d: [0 to 255]D<br><br><i>prefix</i> — Specifies the IPv6 prefix length for the <i>ipv6-address</i> expressed as a decimal integer.<br><br><b>Values</b> 1 to 128<br><br><i>mask</i> — Specifies the eight 16-bit hexadecimal pieces representing bit match criteria.<br><br><b>Values</b> x:x:x:x:x:x:x (eight 16-bit pieces) |

## lsp-exp

|                |                                                                                                                                                            |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>lsp-exp</b> <i>lsp-exp-value</i> [ <b>fc</b> <i>fc-name</i> ] [ <b>priority</b> { <b>low</b>   <b>high</b> }]<br><b>no lsp-exp</b> <i>lsp-exp-value</i> |
| <b>Context</b> | config>qos>sap-ingress                                                                                                                                     |

**Description** This command explicitly sets the forwarding class or subclass enqueueing 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 enqueueing 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 enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing 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.

**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 enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.

A maximum of eight *lsp-exp* rules are allowed on a single policy.

**Values** 0 to 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 subclass-name parameter is optional and used with the *fc-name* parameter to define a preexisting subclass. The *fc-name* and subclass-name parameters must be separated by a period (dot). If subclass-name does not exist in the context of *fc* - name, an error will occur.

**Values**

```
fc: class[.subclass]
 class: be, l2, af, l1, h2, ef, h1, nc
 subclass: 29 characters max
```



**priority** — The priority parameter is used to override the default enqueueing 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 enqueueing 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 enqueueing priority is inherited based on previous rule matches.

**high** — The high parameter is used in conjunction with the priority parameter. Setting the enqueueing parameter to high for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**low** — The low parameter is used in conjunction with the priority parameter. Setting the enqueueing parameter to low for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**Default** no override

## mac-criteria

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mac-criteria</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>The <b>mac-criteria</b> based SAP ingress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.</p> <p>This command is used to enter the node to create or edit policy entries that specify MAC criteria.</p> <p>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.</p> <p>The <b>no</b> form of this command deletes all the entries specified under this node. When mac-criteria entries are removed from a SAP ingress policy, the mac-criteria is removed from all services where that policy is applied.</p> |

## match

|                |                                                                                                         |
|----------------|---------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>match [frame-type {802dot3   802dot2-llc   802dot2-snap   ethernet-II   atm}]</b><br><b>no match</b> |
| <b>Context</b> | config>qos>sap-ingress>mac-criteria>entry                                                               |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>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.</p> <p>If more than one match criteria (within one match statement) are configured, all criteria must be satisfied (AND function) before the action associated with the match will be executed.</p> <p>A <b>match</b> context can consist of multiple match criteria, but multiple <b>match</b> statements cannot be entered per entry.</p> <p>The <b>no</b> form of this command removes the match criteria for the <i>entry-id</i>.</p>                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><b>frame-type</b> — The <b>frame-type</b> keyword configures an Ethernet frame type or an ATM frame type to be used for the MAC filter match criteria.</p> <p><b>Values</b>      802dot3, 802dot2-llc, 802dot2-snap, ethernet_II, atm</p> <p><b>Default</b>      802dot3</p> <p><b>802dot3</b> — Specifies the frame type is Ethernet IEEE 802.3.</p> <p><b>802dot2-llc</b> — Specifies the frame type is Ethernet IEEE 802.2 LLC.</p> <p><b>802dot2-snap</b> — Specifies the frame type is Ethernet IEEE 802.2 SNAP.</p> <p><b>ethernet-II</b> — Specifies the frame type is Ethernet Type II.</p> <p><b>atm</b> — Specifies the frame type as ATM cell. 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. This parameter applies only to the 7750 SR ad 7950 XRS.</p> |

## atm-vci

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>atm-vci</b> <i>vci-value</i><br><b>no atm-vci</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>qos>sap-ingress>mac-criteria>entry>match                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command configures a VCI based filter entry in the SAP ingress QoS policy.</p> <p>This new criterion only takes effect 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/VP RN service has no effect.</p> <p>The user is not allowed to configure a MAC matching criterion other than atm-vci when a MAC criteria filter entry that includes the frame type of atm has been configured.</p> |

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 in which forwarding class queue this packet will be stored. If the user entered a priority value in the priority option, it is ignored as the priority and profile of ATM VLL service packets are 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 forward the packet to the ATM MDA (CMA). The ATM MDA (CMA) stores the individual cells 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.

Values1, 2, 5 to 65535

dot1p

Syntax

dot1p dot1p-value [dot1p-mask]  
no dot1p

Context

config>qos>sap-ingress>mac-criteria>entry>match

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

no dot1p

Parameters

*dot1p-value* — Enter the IEEE 802.1p value in decimal.

Values0 to 7

*dot1p mask* — 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.

**Values** 1 to 7 (decimal)  
**Default** 7

dsap

| <b>Syntax</b>      | <b>dsap</b> <i>dsap-value</i> [ <i>dsap-mask</i> ]<br><b>no dsap</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |              |               |         |         |     |     |             |      |      |        |            |            |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|---------------|---------|---------|-----|-----|-------------|------|------|--------|------------|------------|
| <b>Context</b>     | config>qos>sap-ingress>mac-criteria>entry>match                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |              |               |         |         |     |     |             |      |      |        |            |            |
| <b>Description</b> | <p>Configures an Ethernet 802.2 LLC DSAP value or range for an ingress SAP QoS policy match criterion.</p> <p>This is a 1-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.</p> <p>The snap-pid field, etype field, ssap, and dsap fields are mutually exclusive and cannot be part of the same match criteria.</p> <p>Use the <b>no</b> form of this command to remove the dsap value as the match criterion.</p>                                                                                                                                                                                                      |              |               |         |         |     |     |             |      |      |        |            |            |
| <b>Default</b>     | no dsap                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |              |               |         |         |     |     |             |      |      |        |            |            |
| <b>Parameters</b>  | <p><i>dsap-value</i> — The 8-bit dsap match criteria value in hexadecimal.</p> <p><b>Values</b> 0x00 to 0xFF (hex)</p> <p><i>dsap-mask</i> — This is optional and can be used when specifying a range of dsap values to use as the match criteria.</p> <p>This 8-bit mask can be configured using the following formats:</p> <table><tr><th>Format Style</th><th>Format Syntax</th><th>Example</th></tr><tr><td>Decimal</td><td>DDD</td><td>240</td></tr><tr><td>Hexadecimal</td><td>0xHH</td><td>0xF0</td></tr><tr><td>Binary</td><td>0bBBBBBBBB</td><td>0b11110000</td></tr></table> <p><b>Values</b> 0x00 to 0xFF (hex)</p> <p><b>Default</b> FF</p> | Format Style | Format Syntax | Example | Decimal | DDD | 240 | Hexadecimal | 0xHH | 0xF0 | Binary | 0bBBBBBBBB | 0b11110000 |
| Format Style       | Format Syntax                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Example      |               |         |         |     |     |             |      |      |        |            |            |
| Decimal            | DDD                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 240          |               |         |         |     |     |             |      |      |        |            |            |
| Hexadecimal        | 0xHH                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 0xF0         |               |         |         |     |     |             |      |      |        |            |            |
| Binary             | 0bBBBBBBBB                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 0b11110000   |               |         |         |     |     |             |      |      |        |            |            |

dst-mac

|               |                                                                                      |
|---------------|--------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>dst-mac</b> <i>ieee-address</i> [ <i>ieee-address-mask</i> ]<br><b>no dst-mac</b> |
|---------------|--------------------------------------------------------------------------------------|

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                     |                      |                |         |                  |                 |             |                  |                  |        |               |                |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|----------------------|----------------|---------|------------------|-----------------|-------------|------------------|------------------|--------|---------------|----------------|
| <b>Context</b>      | config>qos>sap-ingress>mac-criteria>entry>match                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                     |                      |                |         |                  |                 |             |                  |                  |        |               |                |
| <b>Description</b>  | <p>Configures a destination MAC address or range to be used as a Service Ingress QoS policy match criterion.</p> <p>The <b>no</b> form of this command removes the destination MAC address as the match criterion.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                     |                      |                |         |                  |                 |             |                  |                  |        |               |                |
| <b>Default</b>      | no dst-mac                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                     |                      |                |         |                  |                 |             |                  |                  |        |               |                |
| <b>Parameters</b>   | <p><i>ieee-address</i> — The MAC address to be used as a match criterion.</p> <p><b>Values</b>        HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit</p> <p><i>ieee-address-mask</i> — A 48-bit mask to match a range of MAC address values.</p> <p>This 48-bit mask can be configured using the following formats:</p> <table><tr><td><b>Format Style</b></td><td><b>Format Syntax</b></td><td><b>Example</b></td></tr><tr><td>Decimal</td><td>DDDDDDDDDDDDDDDD</td><td>281474959933440</td></tr><tr><td>Hexadecimal</td><td>0xHHHHHHHHHHHHHH</td><td>0xFFFFFFFF000000</td></tr><tr><td>Binary</td><td>0bBBBBBBB...B</td><td>0b11110000...B</td></tr></table> <p>All packets with a source MAC OUI value of 00-03-FA, subject to a match condition, should be specified as: 0003FA000000 0xFFFFFFFF000000</p> <p><b>Values</b>        0x0000000000000000 to 0xFFFFFFFFFFFFFFF (hex)</p> <p><b>Default</b>        0xFFFFFFFFFFFFFFF</p> | <b>Format Style</b> | <b>Format Syntax</b> | <b>Example</b> | Decimal | DDDDDDDDDDDDDDDD | 281474959933440 | Hexadecimal | 0xHHHHHHHHHHHHHH | 0xFFFFFFFF000000 | Binary | 0bBBBBBBB...B | 0b11110000...B |
| <b>Format Style</b> | <b>Format Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>Example</b>      |                      |                |         |                  |                 |             |                  |                  |        |               |                |
| Decimal             | DDDDDDDDDDDDDDDD                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 281474959933440     |                      |                |         |                  |                 |             |                  |                  |        |               |                |
| Hexadecimal         | 0xHHHHHHHHHHHHHH                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 0xFFFFFFFF000000    |                      |                |         |                  |                 |             |                  |                  |        |               |                |
| Binary              | 0bBBBBBBB...B                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 0b11110000...B      |                      |                |         |                  |                 |             |                  |                  |        |               |                |

## etype

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>etype</b> <i>etype-value</i></p> <p><b>no etype</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>qos>sap-ingress>mac-criteria>entry>match                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>Configures an Ethernet type II value to be used as a service ingress QoS policy match criterion.</p> <p>The Ethernet type field is a 2-byte field used to identify the protocol carried by the Ethernet frame. For e.g. 0800 is used to identify the IPv4 packets.</p> <p>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.</p> <p>The snap-pid field, etype field, ssap, and dsap fields are mutually exclusive and cannot be part of the same match criteria.</p> |

The no form of this command removes the previously entered etype field as the match criteria.

|                   |                                                                                                                           |
|-------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no etype                                                                                                                  |
| <b>Parameters</b> | <i>etype-value</i> — The Ethernet type II frame Ethertype value to be used as a match criterion expressed in hexadecimal. |
| <b>Values</b>     | 0x0600 to 0xFFFF                                                                                                          |

## inner-tag

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>inner-tag</b> <i>value</i> [ <i>vid-mask</i> ]<br><b>no inner-tag</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>qos>sap-ingress>mac-criteria>entry>match                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command configures the matching of the second tag that is carried transparently through the service. The inner tag on ingress is the second tag on the frame if there are no service delimiting tags. The inner tag is the second tag before any service delimiting tags on egress, but is dependent in the ingress configuration and may be set to 0 even in cases where additional tags are on the frame. This allows matching VLAN tags for explicit filtering or QoS setting when using default or null encapsulations. |

The inner tag is not applicable in ingress on dot1Q SAPs. The inner tag may be populated on egress depending on the ingress SAP type.

On QinQ SAPs of null and default that do not strip tags, the inner-tag will contain the second tag (which is still the second tag carried transparently through the service.) On ingress SAPs that strip any tags, the inner tag will contain 0 even if there are more than two tags on the frame.

The optional vid\_mask is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is ((value and vid-mask) == (tag and vid-mask)). A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6.

For QoS, the VID type cannot be specified on the default QoS policy.

The default vid-mask is set to 4095 for exact match.

## outer-tag

|                |                                                                          |
|----------------|--------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>outer-tag</b> <i>value</i> [ <i>vid-mask</i> ]<br><b>no outer-tag</b> |
| <b>Context</b> | config>qos>sap-ingress>mac-criteria>entry>match                          |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command configures the matching of the first tag that is carried transparently through the service. Service delimiting tags are stripped from the frame and the outer tag on ingress is the first tag after any service delimiting tags. The outer tag is the first tag before any service delimiting tags on egress. This allows matching VLAN tags for explicit filtering or QoS setting when using default or null encapsulations.</p> <p>On dot1Q SAPs, the outer tag is the only tag that can be matched. On dot1Q SAPs with exact match (sap 2/1/1:50), the outer tag will be populated with the next tag that is carried transparently through the service or 0 if there is no additional VLAN tags on the frame.</p> <p>On QinQ SAPs that strip a single service delimiting tag, the outer tag will contain the next tag (which is still the first tag carried transparently through the service.) On SAPs with two service delimiting tags (two tags stripped), the outer-tag will contain 0 even if there are more than two tags on the frame.</p> <p>The optional <i>vid_mask</i> is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is ((value &amp; vid-mask) == (tag &amp; vid-mask)). A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6.</p> <p>For QoS, the VID type cannot be specified on the default QoS policy.</p> <p>The default vid-mask is set to 4095 for exact match.</p> |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## snap-oui

|                    |                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>snap-oui {zero   non-zero}</b><br><b>no snap-oui</b>                                                                                                                                                                                 |
| <b>Context</b>     | config>qos>sap-ingress>mac-criteria>entry>match                                                                                                                                                                                         |
| <b>Description</b> | <p>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.</p> <p>The <b>no</b> form of this command removes the criterion from the match criteria.</p> |
| <b>Default</b>     | no snap-oui                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><b>zero</b> — Specifies to match packets with the 3-byte OUI field in the SNAP-ID set to zero.</p> <p><b>non-zero</b> — Specifies to match packets with the 3-byte OUI field in the SNAP-ID not set to zero.</p>                     |

## snap-pid

|                |                                                 |
|----------------|-------------------------------------------------|
| <b>Syntax</b>  | <b>snap-pid snap-pid</b><br><b>no snap-pid</b>  |
| <b>Context</b> | config>qos>sap-ingress>mac-criteria>entry>match |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>Configures an IEEE 802.3 LLC SNAP Ethernet frame PID value to be used as a service ingress QoS policy match criterion.</p> <p>This is a 2-byte protocol id that is part of the IEEE 802.3 LLC SNAP Ethernet Frame that follows the 3-byte OUI field.</p> <p>The snap-pid field, etype field, ssap, and dsap fields are mutually exclusive and cannot be part of the same match criteria.</p> <p>The <b>snap-pid</b> match criteria is independent of the OUI field within the SNAP header. Two packets with different 3-byte OUI fields, but the same PID field, will both match the same policy entry based on a snap-pid match criteria.</p> <p>The <b>no</b> form of this command removes the snap-pid value as the match criteria.</p> |
| <b>Default</b>     | no snap-pid                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><i>snap-pid</i> — The 2-byte snap-pid value to be used as a match criterion in hexadecimal.</p> <p><b>Values</b>      0x0000 to 0xFFFF</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

src-mac

|                    |                                                                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>src-mac</b> <i>ieee-address</i> [ <i>ieee-address-mask</i> ]<br><b>no src-mac</b>                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>sap-ingress>mac-criteria>entry>match                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command configures a source MAC address or range to be used as a service ingress QoS policy match criterion.</p> <p>The <b>no</b> form of this command removes the source mac as the match criteria.</p>                                                                                  |
| <b>Default</b>     | no src-mac                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>ieee-address</i> — Enter the 48-bit IEEE MAC address to be used as a match criterion.</p> <p><b>Values</b>      HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit</p> <p><i>ieee-address-mask</i> — This 48-bit mask can be configured using the following formats:</p> |

| Format Style | Format Syntax    | Example         |
|--------------|------------------|-----------------|
| Decimal      | DDDDDDDDDDDDDDDD | 281474959933440 |
| Hexadecimal  | 0xHHHHHHHHHHHHH  | 0x0FFFFFF000000 |
| Binary       | 0BBBBBBBB...B    | 0b11110000...B  |



To configure all packets with a source MAC OUI value of 00-03-FA to be subject to a match condition, the entry should be specified as: 003FA000000  
0xFFFFFFFF000000

**Values** 0x0000000000000000 to 0xFFFFFFFFFFFFFFF (hex)

**Default** 0xFFFFFFFFFFFFFFF (hex) (exact match)

## ssap

**Syntax** **ssap** *ssap-value* [*ssap-mask*]  
**no ssap**

**Context** config>qos>sap-ingress>mac-criteria>entry>match

**Description** This command configures an Ethernet 802.2 LLC SSAP value or range for an ingress SAP QoS policy match criterion.

This is a 1-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** no ssap

**Parameters** *ssap-value* — The 8-bit ssap match criteria value in hex.

**Values** 0x00 to 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 |

**Values** 0x00 to 0xFF

## type

**Syntax** **type** *filter-type*  
**no type**

**Context** config>qos>sap-ingress>mac-criteria

---

|                    |                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command sets the mac-criteria type.                                                                                                                                                                                               |
| <b>Default</b>     | type normal                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <i>filter-type</i> — Specifies which type of entries this MAC filter can contain.                                                                                                                                                      |
| <b>Values</b>      | <p><b>normal</b> — Regular match criteria are allowed; ISID match not allowed.</p> <p><b>vid</b> — Configures the VID filter type used to match on ethernet_II frame types. This allows matching VLAN tags for explicit filtering.</p> |

## policer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>policer</b> <i>policer-id</i> [ <b>create</b> ]<br><b>no policer</b> <i>policer-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>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 be defined to have up to 63 policers (numbered 1 through 63) 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 or multiservice sites associated with the policy until a forwarding class is mapped to the policer's ID.</p> |

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.

When 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 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.

When 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 or multiservice sites associated with the QoS policy when 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 to 63

## adaptation-rule

**Syntax** **adaptation-rule** [*pir adaptation-rule*] [*cir adaptation-rule*]  
**no adaptation-rule**

**Context** config>qos>sap-ingress>policer  
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 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 that it can use. 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 adaptation-rule* — When the optional **pir** parameter is specified, the **max**, **min**, or **closest** keyword qualifier must follow.

**Values** **max** — Specifies 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** — Specifies 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 — Specifies 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 adaptation-rule** — When the optional **cir** parameter is specified, the **max**, **min**, or **closest** keyword qualifier must follow.

**Values**

max — Specifies 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 — Specifies 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 — Specifies 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

## adv-config-policy

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>adv-config-policy</b> <i>policy-name</i><br><b>no adv-config-policy</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>qos>sap-ingress>policer<br>config>qos>sap-egress>policer<br>config>qos>sap-ingress>queue<br>config>qos>sap-egress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command specifies the advanced QoS policy. The advanced QoS policy contains only queue and policer child control parameters within a child-control node.</p> <p>When a policy is created, it may be applied to a queue or policer defined within a <b>sap-egress</b> or <b>sap-ingress</b> QoS policy. It may also be applied to a queue or policer defined within an <b>ingress</b> or <b>egress queue-group</b> 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).</p> <p>The <b>no</b> form of this command removes the specified advanced policy.</p> |
| <b>Default</b>     | no adv-config-policy                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>policy-name</i> — The name of the advanced QoS policy.</p> <p><b>Values</b> Valid names consist of any string up to 63 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p>                                                                                                                                                                                                                                                                                                                                                                                       |

## high-prio-only

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>high-prio-only</b> <i>percent-of-mbs</i><br><b>no high-prio-only</b>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>qos>sap-egress>policer<br>config>qos>sap-ingress>policer                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | 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 <b>mbs</b> value defines the policer's high-priority violate threshold, the percentage value defined is applied to the <b>mbs</b> value to derive the bucket's low-priority violate threshold. See the <b>mbs</b> command details for information about which types of traffic are associated with each violate threshold. |
| <b>Parameters</b>  | <i>percent-of-mbs</i> — The <i>percent-of-mbs</i> parameter is required when specifying <b>high-prio-only</b> and is expressed as a percentage.                                                                                                                                                                                                                                                                                                                                                |
| <b>Values</b>      | 0 to 100                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>     | 10                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## percent-rate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>percent-rate</b> <i>pir-percent</i> [ <i>cir cir-percent</i> ]<br><b>no percent-rate</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-ingress>policer<br>config>qos>sap-egress>policer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>The percent-rate command within the SAP ingress and egress QoS policies 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.</p> <p>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.</p> <p>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.</p> <p>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 any time.</p> <p>The <b>no</b> form of this command returns the queue to its default shaping rate and cir rate.</p> |

---

|                   |                                                                                                                                                                                                                                               |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>pir-percent</i> — Specifies the policer's PIR as a percentage of the immediate parent root policer/arbitrator rate or the FP capacity.                                                                                                     |
| <b>Values</b>     | Percentage ranging from 0.01 to 100.00                                                                                                                                                                                                        |
| <b>Default</b>    | 100.00                                                                                                                                                                                                                                        |
|                   | <i>cir-percent</i> — The <b>cir</b> keyword is optional and, when defined, the required <i>cir-percent</i> CIR parameter expresses the policer's CIR as a percentage of the immediate parent root policer/arbitrator rate or the FP capacity. |
| <b>Values</b>     | Percentage ranging from 0.00 to 100.00                                                                                                                                                                                                        |
| <b>Default</b>    | 100.00                                                                                                                                                                                                                                        |

## profile-capped

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] profile-capped</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>qos>sap-ingress>policer<br>config>qos>sap-egress>policer<br>config>qos>qgrps>ingress>queue-group<br>config>qos>qgrps>egr>queue-group                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>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.</p> <p>The profile-capped mode makes two changes:</p> <ul style="list-style-type: none"> <li>• 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.</li> <li>• 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)</li> </ul> |
| <b>Default</b>     | no profile-capped                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## rate

|                |                                                                   |
|----------------|-------------------------------------------------------------------|
| <b>Syntax</b>  | <b>rate {max   <i>pir-rate</i>} [cir {max   <i>cir-rate</i>}]</b> |
| <b>Context</b> | config>qos>sap-ingress>policer<br>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 each packet's size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches its 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 nor 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 kb/s (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 | pir-rate}** — Specifying the keyword **max** or an explicit *pir-rate* 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 *pir-rate* 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 to 2000000000

**cir {*max* | *cir-rate*}** — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword **max** or an explicit *cir-rate* 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 kb/s. The *cir-rate* 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 CIR used is equivalent to **max**.

**Values**      **max** or 0 to 2000000000

## 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 enqueueing 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 enqueueing 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 enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing 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 eight 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 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 HSMDA queue offered stats represent all packets sent to a specific ingress or egress queue regardless of the HSMDA counter override. This results in accurate queue offered stats, while the discard and forwarding stats per queue only represent packets that have not been associated with an exception counter. If the queue discard and forwarding stats are subtracted from the queue offered stats, an approximation of the number of packets handled



by the queue that have been associated with an exception counter may be calculated. This is an approximation due to the possible presence of packets currently in the queue that are not represented by the discard or forwarding stats at the time the stats are collected but had been included in the queue offered stats. This discrepancy is minimized when the stats are collected over time and disappears completely when the queue drains.

The **no** form of this 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 enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing 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 to 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 *subclass-name* parameter is optional and used with the *fc-name* parameter to define a pre-existing subclass. The *fc-name* and *subclass-name* parameters must be separated by a period (.). If *subclass-name* does not exist in the context of *fc-name*, an error will occur. If *subclass-name* is removed using the **no fc fc-name subclass-name force** command, the **default-fc** command will automatically drop the *subclass-name* and only use *fc-name* (the parent forwarding class for the subclass) as the forwarding class.

#### Values

fc: *class[.subclass]*  
*class*: be, l2, af, l1, h2, ef, h1, nc  
*subclass*: 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 enqueueing 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 enqueueing 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 enqueueing priority is inherited based on previous rule matches.

**Values** high, low

**Default** Inherits the priority defined by the default-priority statement.

**high** — This parameter is used in conjunction with the **priority** parameter. Setting the enqueueing parameter to **high** for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**low** — This parameter is used in conjunction with the **priority** parameter. Setting the enqueueing parameter to **low** for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

## queue

**Syntax** `queue queue-id [multipoint] [queue-type] [queue-mode] [create]`  
`no queue queue-id`

**Context** config>qos>sap-ingress

**Description** This command creates the context to configure an ingress SAP QoS policy queue.

Explicit definition of an ingress queue's type is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or best effort 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 best effort 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 queue type 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.

**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 to 32

*queue-type* — The **expedite**, **best-effort**, and **auto-expedite** queue types are mutually exclusive. Each defines the method that the system uses to service the queue from a hardware perspective. A keyword can be specified at the time the queue is created. If an attempt to change the keyword after the queue is initially defined, an error is generated.

**Values** expedite, best-effort, auto-expedite

**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 optional 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 forwarding class unicast traffic is mapped 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 preexisting multipoint queue to edit *queue-id* parameters.

**Default** non-multipoint (unicast queue)

**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

**create** — Keyword creates an ingress SAP QoS policy queue.

## adaptation-rule

|                    |                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>adaptation-rule</b> [ <b>pir</b> <i>adaptation-rule</i> ] [ <b>cir</b> <i>adaptation-rule</i> ]<br><b>no adaptation-rule</b>                                                                                                                                                                  |
| <b>Context</b>     | config>qos>sap-egress>queue                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command defines the method used by the system to derive the operational FIR, CIR, and PIR settings when the queue is provisioned in hardware. For the FIR, CIR, and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint. |

When a specific **adaptation-rule** is removed, the default constraints for **pir**, **cir**, and **fir** apply.

The **no** form of this command removes any explicitly defined constraints used to derive the operational FIR, CIR, and PIR created by the application of the policy.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | adaptation-rule pir closest cir closest fir closest                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b> | <p><b>pir</b> <i>adaptation-rule</i> — Defines the constraints enforced when adapting the queue's PIR. The <b>pir</b> parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the <b>pir</b> parameter is not specified, the default applies.</p> <p><b>Values</b></p> <p><b>max</b> - The <b>max</b> option is mutually exclusive to the <b>min</b> and <b>closest</b> options. When <b>max</b> is specified, the operational rate for the queue will be equal to or less than the requested rate.</p> <p><b>min</b> - The <b>min</b> option is mutually exclusive to the <b>max</b> and <b>closest</b> options. When <b>min</b> is specified, the operational PIR for the queue will be equal to or greater than the requested rate.</p> <p><b>closest</b> - The <b>closest</b> parameter is mutually exclusive to the <b>min</b> and <b>max</b> parameter. When <b>closest</b> is specified, the operational PIR for the queue will be the rate closest to the requested rate.</p> <p><b>cir</b> <i>adaptation-rule</i> — Defines the constraints enforced when adapting the queue's CIR. The <b>cir</b> parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the <b>cir</b> parameter is not specified, the default constraint applies.</p> <p><b>max</b> — Specifies that the operational rate for the queue will be equal to or less than the requested rate.</p> <p><b>min</b> — Specifies that the operational PIR for the queue will be equal to or greater than the requested rate.</p> <p><b>closest</b> — Specifies that the operational PIR for the queue will be the rate closest to the requested rate.</p> |

## adaptation-rule

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>adaptation-rule</b> [ <b>pir</b> <i>adaptation-rule</i> ] [ <b>cir</b> <i>adaptation-rule</i> ] [ <b>fir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }]<br><b>no adaptation-rule</b>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>sap-ingress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command defines the method used by the system to derive the operational FIR, CIR and PIR settings when the queue is provisioned in hardware. For the FIR, CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.</p> <p>When a specific <b>adaptation-rule</b> is removed, the default constraints for <b>pir</b>, <b>cir</b> and <b>fir</b> apply.</p> <p>The <b>no</b> form of this command removes any explicitly defined constraints used to derive the operational FIR, CIR and PIR created by the application of the policy.</p> |
| <b>Default</b>     | adaptation-rule pir closest cir closest fir closest                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <b>pir</b> <i>adaptation-rule</i> — Defines the constraints enforced when adapting the queue's PIR. The <b>pir</b> parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the <b>pir</b> parameter is not specified, the default applies.                                                                                                            |
|                   | <b>cir</b> <i>adaptation-rule</i> — Defines the constraints enforced when adapting the queue's CIR. The <b>cir</b> parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the <b>cir</b> parameter is not specified, the default constraint applies.                                                                                                 |
|                   | <b>fir</b> — Defines the constraints enforced when adapting the queue's FIR. The <b>fir</b> parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the <b>fir</b> parameter is not specified, the default constraint applies. FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware. |
|                   | <b>max</b> — Specifies that the operational rate for the queue will be equal to or less than the requested rate.                                                                                                                                                                                                                                                                                                    |
|                   | <b>min</b> — Specifies that the operational rate for the queue will be equal to or greater than the requested rate.                                                                                                                                                                                                                                                                                                 |
|                   | <b>closest</b> — Specifies that the operational rate for the queue will be the rate closest to the requested rate.                                                                                                                                                                                                                                                                                                  |

## burst-limit

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>burst-limit</b> <i>size</i> [ <b>bytes</b>   <b>kilobytes</b> ]<br><b>no burst-limit</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>qos>sap-ingress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>The <b>queue burst-limit</b> 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.</p> <p>The <b>burst-limit</b> 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.</p> <p>The <b>no</b> 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 <b>queue-override queue</b> 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.</p> |
| <b>Default</b>     | no burst-limit                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><b>size</b> — Specifies 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.</p> <p><b>Values</b> 1 to 13,671 kbytes or 14,000,000 bytes</p> <p><b>Default</b> No default for size; use the default keyword to specify default burst limit.</p> <p><b>bytes</b> — Specifies that the value given for size must be interpreted as the burst limit in bytes.</p> <p><b>kilobytes</b> — Specifies that the value given for size must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.</p> |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## cbs

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>cbs</b> <i>size-in-kbytes</i></p> <p><b>no cbs</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | <p>config&gt;qos&gt;sap-egress&gt;queue</p> <p>config&gt;qos&gt;sap-ingress&gt;queue</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>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.</p> <p>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.</p> <p>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.</p> <p>The <b>no</b> form of this command returns the CBS size to the default value.</p> |
| <b>Default</b>     | cbs default                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><i>size-in-kbytes</i> — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes 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) The CBS maximum value used is constrained by the pool size in which the queue exists.</p> <p><b>Values</b> 0 to 1048576 or default</p> <p>Minimum configurable non-zero value: 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4</p> <p>Minimum non-zero default value: maximum of 10 ms of CIR, or 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4</p> |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## cir-non-profiling

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] cir-non-profiling</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>sap-ingress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command prevents the modification of the profile of a packet depending on the queue rate compared to its configured CIR. The CIR continues to be used to affect the scheduling priority of a queue. The <b>cir-non-profiling</b> command and the <b>queue police</b> command are mutually exclusive.</p> <p>The <b>cir-non-profiling</b> command is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.</p> <p>The <b>cir-non-profiling</b> command should not be configured under a SAP ingress QoS policy queue associated with a LAG which spans FP4-based and FP2- or FP3-based hardware as the resulting operation could be different depending on which hardware type the traffic ingresses.</p> |

## drop-tail

|                    |                                                                          |
|--------------------|--------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>drop-tail</b>                                                         |
| <b>Context</b>     | config>qos>sap-egress>queue<br>config>qos>sap-ingress>queue              |
| <b>Description</b> | This command enters the context to configure queue drop tail parameters. |

## low

|                |                                       |
|----------------|---------------------------------------|
| <b>Syntax</b>  | <b>low</b>                            |
| <b>Context</b> | config>qos>sap-egress>queue>drop-tail |



config>qos>sap-ingress>queue>drop-tail

**Description** This command enters the context to configure the queue low drop-tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets will not be accepted into the queue and will be discarded.

## percent-reduction-from-mbs

**Syntax** **percent-reduction-from-mbs** *percent*  
**no percent-reduction-from-mbs**

**Context** config>qos>sap-ingress>queue>drop-tail>low

**Description** This command configures the ingress SAP low drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and this percentage is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value and will be discarded.

**Default** 10%

**Parameters** *percent* — Specifies the percentage reduction from the MBS for a queue drop tail.

**Values** 0 to 100, default

## mbs

**Syntax** **mbs** {*size* [**bytes** | **kilobytes**] | **default**}  
**no mbs**

**Context** config>qos>sap-egress>hsmda-queues>queue  
config>qos>sap-egress>queue  
config>qos>sap-ingress>queue

**Description** This command configures the maximum number of buffers allowed for a specific queue. The value is given in bytes or kilobytes and overrides the default value for the context.

The **no** form of this command returns the policer to its default MBS.

**Default** no mbs

**Parameters** *size* — The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes.

**Default** **kilobyte**

**default** — Sets the MBS to its default value.

**bytes** — Specifies that the value given for size is interpreted as the queue's MBS value given in bytes.

**Values** 0 to 2688000

**kilobytes** — Specifies the value is interpreted as the queue's MBS value given in kilobytes.

**Values** 0 to 2625

**Default** kilobytes

## packet-byte-offset

**Syntax** **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}  
**no packet-byte-offset**

**Context** config>qos>sap-ingress>queue

**Description** This command modifies the size of each packet handled by the queue by adding or subtracting the specified number of bytes. The actual packet size is not modified; only the size used to determine the ingress scheduling and profiling is changed. The **packet-byte-offset** command is an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the scheduling and profiling throughput is affected by the offset as well as the statistics (accounting) associated with the queue. The **packet-byte-offset** does not apply to drop statistics, received valid statistics, or the offered managed and unmanaged statistics used by Ingress Multicast Path Management.

The **no** form of this command removes per-packet size modifications from the queue.

**Parameters** *add-bytes* — Specifies the number of bytes added to the size of each packet associated with the queue for scheduling, profiling, and accounting purposes. From the queue's perspective, the packet size is increased by the amount specified.

**Values** 0 to 30, in increments of 2

*sub-bytes* — Specifies the number of bytes subtracted from the size of each packet associated with the queue for scheduling, profiling, and accounting purposes. From the queue's perspective, the packet size is reduced by the amount specified. The minimum resulting packet size used by the system is 1 byte.

**Values** Values 0 to 64, in increments of 2

## 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, policers (at egress only), 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 multiservice 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 SAP that the queue belongs to also depicts 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 and automatically returns to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying the weight parameter, the default is a weight of 1.

The **no** form of this 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. When 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 that 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, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

*weight* — *weight* defines the relative weight of this queue in comparison to other child schedulers, policers, and queues, while vying for bandwidth on the parent *scheduler-name*. Any queues, policers, or schedulers defined as weighted receive no parental bandwidth until all policers, queues, and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

All **weight** values from all weighted active queues, policers, 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, policer, or scheduler. A weight is considered to be active when the pertaining queue, policer, or scheduler has not reached its maximum rate and still has packets to transmit. All child queues, policers, and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all non-zero weighted queues, policers, and schedulers at that level are operating at the maximum bandwidth or are idle.

**Values** 0 to 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*. Queues or schedulers will not receive parental bandwidth until all queues, policers, 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 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 relation to their relative weights.

**Values** 1 to 8

**Default** 1

*cir-weight* — Defines the weight that 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 to 100

**Default** 1

*cir-level* — Defines the port priority that 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 to 8 (8 is the highest priority)

**Default** 0

## percent-rate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>percent-rate</b> <i>pir-percent</i> [ <i>cir cir-percent</i> ] [ <i>fir fir-percent</i> ] [{ <b>port-limit</b>   <b>local-limit</b> }]<br><b>percent-rate</b> <i>pir-percent</i> <b>police</b> [{ <b>port-limit</b>   <b>local-limit</b> }]<br><b>no percent-rate</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>qos>sap-ingress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command configures a queue's PIR and CIR as a percentage of the ingress port line rate or as a percentage of its parent scheduler rate. When the rates are expressed as a <b>port-limit</b>, 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 Gb and a 10 Gb Ethernet port, the queue's rates will be 10 times greater on the 10 Gb 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.</p> <p>If the port's speed changes after the queue is created, the queue's PIR, CIR, and FIR rates will be recalculated based on the defined percentage value.</p> <p>When the rates are expressed as a <b>local-limit</b>, 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, CIR, and FIR will be recalculated based on the defined percentage value.</p> <p>The <b>rate</b> and <b>percent-rate</b> commands override one another. If the current rate for a queue is defined using the <b>percent-rate</b> command and the <b>rate</b> command is executed, the <b>percent-rate</b> values are deleted. Similarly, the <b>percent-rate</b> command causes any <b>rate</b> command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at any time.</p> <p>Queue rate overrides can only be specified in the form configured in the QoS policy (for example, a SAP override can only be specified as a <b>percent-rate</b> if the associated QoS policy was also defined as percent-rate). Likewise, a SAP override can only be specified as a <b>rate</b> (kb/s) if the associated QoS policy was also defined as a rate. Queue-overrides are relative to the limit type specified in the QoS policy.</p> <p>When no <b>percent-rate</b> is defined within a SAP ingress queue-override, the queue uses the defined shaping rate, CIR, and FIR within the SAP ingress QoS policy associated with the queue.</p> <p>The <b>no</b> form of this command returns the queue to its default shaping rate, CIR, and FIR.</p> |
| <b>Parameters</b>  | <p><i>pir-percent</i> — Specifies the queue's PIR as a percentage dependent on the use of the <b>port-limit</b> or <b>local-limit</b>.</p> <p><b>Values</b>      0.01 to 100.00</p> <p><b>Default</b>      100.00</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

*cir-percent* — Specifies the queue's CIR as a percentage dependent on the use of the **port-limit** or **local-limit**.

**Values** 0.00 to 100.00

**Default** 100.00

*fir-percent* — Specifies the queue's FIR as a percentage dependent on the use of the **port-limit** or **local-limit**. FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

**Values** 0.00 to 100.00

**Default** 100.00

**police** — Specifies that traffic feeding into the physical queue instance above the specified PIR rate will be dropped. When the **police** keyword is defined, only the PIR rate may be overridden. The **police** keyword is only applicable to SAP ingress.

**port-limit** — Specifies that the configured PIR, CIR, and FIR percentages are relative to the rate of the port (including the ingress-rate setting) to which the queue is attached. The **port-limit** is the default.

**local-limit** — Specifies that the configured PIR, CIR, and FIR percentages are relative to the queue's parent scheduler rate. If there is no parent scheduler rate, or its rate is **max**, the **port-limit** is used.

## pool

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>pool</b> <i>pool-name</i><br><b>no pool</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>qos>sap-ingress>queue<br>config>qos>sap-egress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command allows the queue to receive buffers from an explicit buffer pool instead of the default buffer pool. The specified <i>pool-name</i> 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. Note that named pools are not supported on the 7950 XSR, 7750 SRc4/12, 7750 SR-a4/a8, or 7750 SR-1e/2e/3e.</p> <p>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 specified <i>pool-name</i> does not exist on the XMA or MDA, the queue will be treated as pool orphaned and will be mapped to the appropriate default pool. When the pool comes into existence on the MDA or port, the queue will be mapped to the new pool.</p> |

When 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 on the **queue** command output using the **pool** keyword, if it exists.

The **no** form of this command is used to remove a named pool association for the queue. When the pool name is removed, the queue will be placed on the appropriate default pool.

**Default** no pool

**Parameters** *pool-name* — Specifies a named pool where the policy will be applied. Each queue created within the system is tied to a physical port. The specified name must be an ASCII name string up to 16 characters.

## rate

**Syntax** **rate** *pir-rate* [**cir** *cir-rate*] [**fir** *fir-rate*]  
**rate** *pir-rate* **police**  
**no rate**

**Context** config>qos>sap-ingress>queue

**Description** This command defines the administrative Peak Information Rate (PIR), the administrative Committed Information Rate (CIR), and the administrative Fair Information Rate (FIR) 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, unless **cir-non-profiling** is configured. In-profile, then out-of-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 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 FIR defines an additional rate at which the system prioritizes the queue over other queues competing for the same bandwidth above that used by the CIR.

The **rate** command can be executed at any time, altering the PIR, CIR, and FIR for all queues created through the association of the SAP ingress QoS policy with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0, 0).

**Default**     rate max cir 0 fir 0

**Parameters**     *pir-rate* — Defines the administrative PIR, 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 the value must be given as a positive integer.

The actual PIR is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values**     1 to 200000000 kb/s, **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 the value must be given as a positive integer.

The actual CIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values**     0 to 200000000 kb/s, **max**

**Default**     0

*fir-rate* — The **fir** parameter overrides the default administrative FIR used by the queue. When the **rate** command is executed, an FIR setting is optional. When the **rate** command has not been executed or the **fir** parameter is not explicitly specified, the default FIR (0) is assumed.

Fractional values are not allowed and the value must be given as a positive integer.

The actual FIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

**Values**     0 to 200000000 kb/s, **max**

**Default**     0

**police** — Specifies that traffic feeding into the queue instance above the specified PIR rate will be dropped. When the **police** keyword is defined, only the PIR rate may be overridden.

## scope

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



---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>qos>sap-ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures the Service Ingress QoS policy scope as exclusive or template.</p> <p>The policy's scope cannot be changed if the policy is applied to a service.</p> <p>The <b>no</b> form of this command sets the scope of the policy to the default of <b>template</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>     | scope template                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><b>exclusive</b> — 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.</p> <p>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.</p> <p><b>template</b> — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs on the router.</p> <p>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.</p> |

## sub-insert-shared-pccrule

|                    |                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sub-insert-shared-pccrule start-entry entry-id count count</b><br><b>no sub-insert-shared-pccrule</b>                                                                                                                                                                                                                            |
| <b>Context</b>     | config>qos>sap-egress<br>config>qos>sap-ingress                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command defines the range of filter and QoS policy entries that are reserved for shared entries received in Flow-Information AVP via Gx interface (PCC rules – Policy and Charging Control).</p> <p>The <b>no</b> form of this command disables the insertion, which will result in a failure of PCC rule installation.</p> |
| <b>Default</b>     | no sub-insert-shared-pccrule                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>entry-id</i> — Specifies the lowest entry in the range.</p> <p><b>Values</b> 1 to 65535</p> <p><i>count</i> — Specifies the number of entries in the range.</p> <p><b>Values</b> 1 to 65535</p>                                                                                                                               |

### 6.6.2.3 Service Egress QoS Policy Commands

#### sap-egress

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sap-egress</b> { <i>policy-id</i>   <i>policy-name</i> } [ <b>create</b> ] [ <b>name</b> <i>name</i> ]<br><b>no sap-egress</b> { <i>policy-id</i>   <i>policy-name</i> }                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>qos                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command is used to create or edit a Service Egress QoS policy. The egress policy defines the SLA for service packets as they egress on the SAP.</p> <p>Policies 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.</p> <p>Sap-egress policies determine queue mappings based on ingress DSCP, IP precedence, dot1p, and IPv4 or IPv6 match criteria. Multiple queues can be created per forwarding class and each queue can have different CIR or PIR parameters.</p> <p>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 DSCP, IP precedence, or dot1p value can optionally be specified.</p> <p>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 default sap-egress policy cannot be modified or deleted.</p> <p>By default, all forwarding classes map to queue 1.</p> <p>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, then written over the original policy-id. Use the <b>config qos copy</b> command to maintain policies in this manner.</p> <p>The <b>no</b> form of this command 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.</p> |
| <b>Parameters</b>  | <p><i>policy-id</i> — The <i>policy-id</i> uniquely identifies the policy on the router.</p> <p><b>Values</b> 1 to 65535</p> <p><i>policy-name</i> — The <i>policy-name</i> uniquely identifies the policy.</p> <p><b>Values</b> 64 characters maximum.</p> <p><b>create</b> — Required parameter when creating a SAP QoS egress policy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

*name* — Configures an optional policy name, up to 64 characters in length, which adds a name identifier to a given policy to then use that policy name in configuration references as well as display and use policy names in show commands throughout the system. This helps the service provider or administrator to identify and manage sap-egress policies within the SR OS platforms.

All sap-egress policies are required to assign a policy ID to initially create a policy. However, either the policy ID or the policy name can be used to identify and reference a given policy once it is initially created

If a name is not specified at creation time, then SR OS assigns a string version of the *policy-id* as the name.

**Values** 64 characters maximum

## dot1p

**Syntax** **dot1p** *dot1p-value* [**fc** *fc-name*] [**profile** {**in** | **out** | **use-de** | **exceed** | **inplus**}]  
**no dot1p** *dot1p-value*

**Context** config>qos>sap-egress

**Description** This command defines a specific dot1p value that must be matched to perform the associated reclassification actions. If an egress packet on the SAP matches the specified dot1p value, the forwarding class or profile may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions.

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 QoS behavior.

The reclassification actions from a dot1p reclassification rule may be overridden by a DSCP, IP precedence, or IP flow matching 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 a DSCP, IP precedence, IPv6 criteria, or IP criteria match occurs after the dot1p 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 dot1p 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 DSCP, IP precedence, IPv6 criteria, or IP criteria match occurs after the dot1p 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 dot1p match will be used.

The **no** form of this command removes the reclassification rule from the SAP egress QoS 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 enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.  
A maximum of eight dot1p rules are allowed on a single policy.
- Values** 0 to 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
- profile {in | out | use-de | exceed | inplus}** — The profile reclassification action is optional. When specified, packets matching the dot1p value will be explicitly reclassified to the profile specified regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by a DSCP, IP precedence, IPv6 criteria, or IP criteria reclassification match. To remove the profile reclassification action for the specified dotp1 value, the **dot1p** command must be re-executed without the profile reclassification action defined.
- Values**
- in** — Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.
  - out** — Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.
  - use-de** — Specifies that the DE bit is used to determine the profile of the packet (in-profile is used when DE = 0 and out-of-profile is used when DE = 1).
  - exceed** — Specifies that any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.
  - inplus** — Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

## dscp

- Syntax** **dscp** *dscp-name* [*dscp-name*] [**fc** *fc-name*] [**hsmdda-counter-override** *counter-id*] [**profile** {in | out | exceed | inplus}]  
**no dscp** *dscp-name* [*dscp-name*]
- 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 eight 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 eight names. Entering more than eight DSCP names with the same parameters (**fc**, **hsmdda-counter-override**, **profile**) will result in multiple lists being created. Conversely, multiple lists with the same parameters (**fc**, **hsmdda-counter-override**, **profile**) are merged and the lists repacked to a maximum of eight 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, a new list will be created for the subset. 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 HSMDDA 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 HSMDDA 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 **hsmdda-counter-override** keyword is optional. When specified, and the egress SAP is created on an HSMDDA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMDDA uses each queue's default queue counters for packets mapped to the queue. The **hsmdda-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 queue's discard or forwarding counters; instead, the exception discard and forwarding counters will be used.

The DSCP-based counter override decision may be overridden by an IP criteria reclassification rule match if the higher priority classification rule has an **hsmda-counter-override** action defined.

The **no** form of this 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, 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 eight 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 following command:

```
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-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** — 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 the specified profile regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by an IPv6 criteria or 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** — Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.
- out** — Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.
- exceed** — Specifies that when **exceed** is specified, any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.
- inplus** — Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

## stat-mode

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>stat-mode {no-stats   minimal   offered-profile-no-cir   offered-profile-cir   offered-total-cir   offered-limited-capped-cir   offered-profile-capped-cir}</b><br><b>no stat mode</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-egress>dyn-policer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | The <b>sap-egress</b> QoS policy's policer <b>stat-mode</b> command is used to configure the forwarding plane counters that allow offered, forwarded, and dropped 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, out-of-profile, and exceed-profile due to egress profile overrides) 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, for example, will not be configured with a CIR profiling rate and not all policers will receive explicitly reprofiled offered packets. The <b>stat-mode</b> 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 that prevents any packet accounting, the use of the policer's **parent** command requires that the policer's **stat-mode** 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. When 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. The total, allocated, and free statistics can be viewed by using the **tools dump resource-usage card fp** 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 egress policer stat-modes are described in [Table 39](#).



**Table 39 Egress Policer Stat-mode Summary**

| Stat Mode                  | Stat Resources | Traffic Counters (Packet/Octets)                                                                                                  |                                                      | Comments                                                                                                                                                                                        |
|----------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                            |                | Offered                                                                                                                           | Dropped/<br>Forwarded                                |                                                                                                                                                                                                 |
| no-stats                   | 0              | None                                                                                                                              | None                                                 |                                                                                                                                                                                                 |
| minimal                    | 1              | Single counter entering policer                                                                                                   | Single counter for dropped/forwarded exiting policer |                                                                                                                                                                                                 |
| offered-profile-no-cir     | 2              | In or out entering policer                                                                                                        | In/out entering policer                              | Intended for when the policer does not change the profile of packets. Includes only in-profile and out-of-profile.                                                                              |
| offered-profile-cir        | 4              | In, out, or uncolored (which corresponds to hard in-profile, hard out-of-profile, or soft in- or out-of-profile) entering policer | In/out exiting policer                               | Intended for when the policer can change the profile of packets to in-profile and out-of-profile.                                                                                               |
| offered-total-cir          | 2              | Single counter entering policer                                                                                                   | In/out exiting policer                               |                                                                                                                                                                                                 |
| offered-limited-capped-cir | 4              | In or out entering policer                                                                                                        | In/out exiting policer                               | Intended for when the policer has <b>profile-capped</b> configured. The information is limited compared to <b>offered-profile-capped-cir</b> with the benefit of using one less stat resources. |
| offered-profile-capped-cir | 5              | In, out, or uncolored (which corresponds to hard in-profile, hard out-of-profile, or soft in- or out-of-profile) entering policer | In/out exiting policer                               | Intended for when the policer has <b>profile-capped</b> configured                                                                                                                              |

When a policer is created within the policy, the default setting for **stat-mode** 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, no statistics are generated.

**minimal** — Counter resource allocation: 1

This stat-mode provides the minimal accounting resource usage and counter information, and includes only the total offered, dropped and forwarded packet and octet counters for traffic entering (offered) and exiting (dropped/forwarded) the policer.

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 and do not count different profile output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate or using exceed PIR.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 40](#).

**Table 40**    **Egress Accounting Statistics Collected in minimal stat-mode**

| Show Output | Accounting Stats Collected |                     |
|-------------|----------------------------|---------------------|
|             | Field                      | Field Description   |
| Off. All    | apo                        | AllPacketsOffered   |
|             | aoo                        | AllOctetsOffered    |
| Dro. All    | apd                        | AllPacketsDropped   |
|             | aod                        | AllOctetsDropped    |
| For. All    | apf                        | AllPacketsForwarded |
|             | aof                        | AllOctetsForwarded  |

**offered-profile-no-cir** — Counter resource allocation: 2

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering the policer. inplus-profile traffic is counted with the in-profile counters and exceed-profile traffic is counted with the out-of-profile counters.

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, dropped, and forwarded stats are required from the egress policer, but a CIR or **enable-exceed-pir** 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 or using **enable-exceed-pir**.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 41](#).

**Table 41 Egress Accounting Statistics Collected in offered-profile-no-cir stat-mode**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. InProf  | ipo                        | InProfilePacketsOffered      |
|              | ioo                        | InProfileOctetsOffered       |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered   |
|              | ooo                        | OutOfProfileOctetsOffered    |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-profile-cir** — Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is made up of traffic that is in-plus-profile, in-profile, out-of-profile, exceed-profile, soft in-profile, and soft out-of-profile. The offered counters cover traffic reclassified to in-profile (which includes traffic reclassified to in-plus-profile), traffic reclassified to out-of-profile (which includes traffic reclassified to exceed-profile) and traffic which has not been reclassified at egress (Uncolor). In the dropped and forwarded counters, in-plus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

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 profile-based offered, dropped and forwarded 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.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 42](#).

**Table 42 Egress Accounting Statistics Collected in offered-profile-cir stat-mode**

| Show Output  | Accounting Stats Collected |                            |
|--------------|----------------------------|----------------------------|
|              | Field                      | Field Description          |
| Off. InProf  | ipo                        | InProfilePacketsOffered    |
|              | ioo                        | InProfileOctetsOffered     |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered |
|              | ooo                        | OutOfProfileOctetsOffered  |
| Off. Uncolor | ucp                        | UncoloredPacketsOffered    |
|              | uco                        | UncoloredOctetsOffered     |
| Dro. InProf  | ipd                        | InProfilePacketsDropped    |
|              | iod                        | InProfileOctetsDropped     |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped |
|              | ood                        | OutOfProfileOctetsDropped  |
| For. InProf  | ipf                        | InProfilePacketsForwarded  |
|              | iof                        | InProfileOctetsForwarded   |

**Table 42 Egress Accounting Statistics Collected in offered-profile-cir stat-mode (Continued)**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-total-cir** — Counter resource allocation: 2

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

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 untrusted 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 untrusted packets.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 43](#).

**Table 43 Egress Accounting Statistics Collected in offered-total-cir stat-mode**

| Show Output  | Accounting Stats Collected |                            |
|--------------|----------------------------|----------------------------|
|              | Field                      | Field Description          |
| Off. All     | apo                        | AllPacketsOffered          |
|              | aoo                        | AllOctetsOffered           |
| Dro. InProf  | ipd                        | InProfilePacketsDropped    |
|              | iod                        | InProfileOctetsDropped     |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped |
|              | ood                        | OutOfProfileOctetsDropped  |

**Table 43** Egress Accounting Statistics Collected in offered-total-cir stat-mode (Continued)

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-limited-capped-cir** — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is made up of traffic that is inplus-profile, in-profile, out-of-profile, exceed-profile, soft in-profile, and soft out-of-profile. The offered counters cover in-profile traffic (which includes traffic reclassified to inplus-profile) and out-of-profile traffic (which includes traffic reclassified to exceed-profile). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

**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 **soft in-profile** with **profile in** and **soft-out-of-profile** with **profile out** and eliminates the offered-undefined statistic.

The impact of using **offered-limited-capped-cir** stat-mode while profile-capped mode is disabled is 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.

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 44](#).

**Table 44 Egress Accounting Statistics Collected in offered-limited-capped-cir stat-mode**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. InProf  | ipo                        | InProfilePacketsOffered      |
|              | ioo                        | InProfileOctetsOffered       |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered   |
|              | ooo                        | OutOfProfileOctetsOffered    |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-profile-capped-cir** — Counter resource allocation: 5

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is made up of traffic that is inplus-profile, in-profile, out-of-profile, exceed-profile, soft in-profile, and soft out-of-profile. The offered counters cover traffic reclassified to in-profile (which includes traffic reclassified to inplus-profile), traffic reclassified to out-of-profile (which includes traffic reclassified to exceed-profile) and traffic that has not been reclassified at egress (uncolored). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

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 inplus**, **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 is 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 (hard in-profile) instead of offered-undefined (uncolored).

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 45](#).

**Table 45 Egress Accounting Statistics Collected in offered-profile-capped-cir stat-mode**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. InProf  | ipo                        | InProfilePacketsOffered      |
|              | ioo                        | InProfileOctetsOffered       |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered   |
|              | ooo                        | OutOfProfileOctetsOffered    |
| Off. Uncolor | ucp                        | UncoloredPacketsOffered      |
|              | uco                        | UncoloredOctetsOffered       |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

## 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** no ethernet-ctag

## fc

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

**Parameters** *fc-name* — This parameter specifies that 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

## de-mark

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

**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 inplus-profile/in-profile or out-of-profile/exceed-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 inplus-profile or in-profile packets, 1 for out-of-profile or exceed-profile packets. For more information, refer to the *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; otherwise, set to zero.

If the DE value is included in the command line, this value is to be used for all the packets of this forwarding class regardless of their 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.

**Parameters** *de-value* — Specifies the DE marking value.

**Values** 0 or 1

## de-mark-inner

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

**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 inplus-profile or in-profile packets, 1 for out-of-profile or exceed-profile packets. For more information, refer to the *IEEE 802.1ad-2005 standard*.

If the DE value is included in the command line, this value is used for all the inner tags of packets of this forwarding class, regardless of their 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. 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).

If no de-mark commands are used, the DE bit is preserved if an ingress inner tag exists; otherwise, set to 0.

Remarking the inner DE bit is not supported based on the profile result of egress policing.

**Parameters** *de-value* — Specifies the DE marking value.

**Values** 0 or 1

## de-mark-outer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>de-mark-outer</b> [ <b>force</b> <i>de-value</i> ]<br><b>no de-mark-outer</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>qos>sap-egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>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, out, or exceed-profile status of the packet.</p> <p>If no DE value is present, the default values are used for the marking of the DE bit; for example, 0 for inplus-profile/in-profile packets, 1 for out-of-profile/exceed-profile packets. For more information, refer to the <i>IEEE 802.1ad-2005 standard</i>.</p> <p>If the DE value is included in the command line, this value is used for all the outer or single tags of packets of this forwarding class, regardless of their profile status.</p> <p>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.</p> <p>This command takes precedence over the <b>de-mark</b> command if both are specified in the same policy and over the default action.</p> <p>If no de-mark commands are used, the DE bit is preserved if an ingress outer or single tag exists; otherwise, set to 0.</p> |
| <b>Parameters</b>  | <p><i>de-value</i> — Specifies the DE marking value.</p> <p><b>Values</b>      0 or 1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## dot1p

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dot1p</b> { <i>dot1p-value</i>   <b>in-profile</b> <i>dot1p-value</i> <b>out-profile</b> <i>dot1p-value</i> [ <b>exceed-profile</b> <i>dot1p-value</i> ]}<br><b>no dot1p</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>qos>sap-egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command explicitly defines the egress IEEE 802.1p (dot1p) bits marking for <i>fc-name</i>. When the marking is set, all packets of <i>fc-name</i> that have either an IEEE 802.1q or IEEE 802.1p encapsulation use the explicitly defined <i>dot1p-value</i>. If the egress packets for <i>fc-name</i> are not IEEE 802.1q or IEEE 802.1p encapsulated, the dot1p command has no effect.</p> <p>The optional <b>in-profile</b> <i>dot1p-value</i> <b>out-profile</b> <i>dot1p-value</i> [<b>exceed-profile</b> <i>dot1p-value</i>] parameters added to the existing <b>dot1p</b> command adds the capability to mark on an egress SAP the in, out, and exceed-profile status via a certain dot1p combination, similarly with the DE options. All inplus-profile traffic is marked with the same value as in-profile traffic.</p> |

When the **in-profile** keyword is added, the **out-profile** keyword must be specified; however, **exceed-profile** is optional. If the optional **exceed-profile dot1p-value** is not included, any exceed-profile traffic will be marked with the same dot1p value as configured for the out-of-profile traffic.

The command with the additional structure may be used on the SAP when the internal in, out, and exceed-profile status needs to be communicated to an access network or customer device that does not support the DE bit.

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, as it did before. The two versions of the command are mutually exclusive.

The in-profile or out-of-profile/exceed-profile status may be indicated via the DE bit setting if the **de-mark** command is used. The DE value used for exceed-profile is the same as that used for out-of-profile.

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 this command sets the IEEE 802.1p or IEEE 802.1q priority bits to 0.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no dot1p                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b> | <p><b>in-profile dot1p-value</b> — Specifies the 802.1p value to set for in-profile frames in this forwarding class.</p> <p><b>Values</b> 0 to 7</p> <p><b>out-profile dot1p-value</b> — Specifies the 802.1p value to set for out-profile frames in this forwarding class.</p> <p><b>Values</b> 0 to 7</p> <p><b>exceed-profile dot1p-value</b> — Specifies the 802.1p value to set for exceed-profile frames in this forwarding class.</p> <p><b>Values</b> 0 to 7</p> |

## dot1p-inner

|                |                                                                                                            |
|----------------|------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>dot1p-inner {dot1p-value   in-profile dot1p-value out-profile dot1p-value}</b><br><b>no dot1p-inner</b> |
| <b>Context</b> | config>qos>sap-egress>fc                                                                                   |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command explicitly defines the egress inner VLAN tag IEEE 802.1p (dot1p) bits marking for <i>fc-name</i>. When the marking is set, all packets of <i>fc-name</i> that have either an inner IEEE 802.1q or IEEE 802.1p encapsulation on a qinq SAP will use the explicitly defined <i>dot1p-value</i>. If the egress packets for <i>fc-name</i> are not IEEE 802.1q or IEEE 802.1p qinq encapsulated, this command has no effect.</p> <p>The optional <b>in-profile</b>   <i>dot1p-value</i> <b>out-profile</b> <i>dot1p-value</i> parameters on the <b>dot1p-inner</b> command add the capability to mark the in-profile 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-profile and out-of-profile status needs to be communicated to an access network or customer device that does not support the DE bit. When the in-profile keyword is added, the rest of the structure must be specified. All inplus-profile traffic is marked with the same value as in-profile traffic and all exceed-profile traffic is marked with the same value as out-of-profile traffic.</p> <p>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 inplus/in or out/exceed-profile status may be indicated via the setting of the DE bit setting if the <b>de-mark(-inner)</b> command is used.</p> <p>The two versions of the command (with and without parameters) are mutually exclusive.</p> <p>This command takes precedence over the <b>dot1p</b> command if both are specified in the same policy, and over the default action where the marking is taken from a packet received at ingress.</p> <p>The configuration of <b>qinq-mark-top-only</b> under the SAP egress takes precedence over the use of the <b>dot1p-inner</b> in the policy; that is, the inner VLAN tag is not remarked when <b>qinq-mark-top-only</b> 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).</p> <p>The <b>no</b> form of the command sets the inner IEEE 802.1p or IEEE 802.1q priority bits to 0.</p> <p>Remarking the inner dot1p is not supported based on the profile result of egress policing.</p> |
| <b>Default</b>     | no dot1p-inner                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>dot1p-value</i> — Specifies the 802.1p value to set for in-profile frames in this forwarding class.</p> <p><b>Values</b> 0 to 7</p> <p><b>in-profile</b> <i>dot1p-value</i> — Specifies the 802.1p value to set for in-profile frames in this forwarding class.</p> <p><b>Values</b> 0 to 7</p> <p><b>out-profile</b> <i>dot1p-value</i> — Specifies the 802.1p value to set for out-of-profile frames in this forwarding class.</p> <p><b>Values</b> 0 to 7</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

---

dot1p-outer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dot1p-outer</b> { <i>dot1p-value</i>   <b>in-profile</b> <i>dot1p-value</i> <b>out-profile</b> <i>dot1p-value</i> [ <b>exceed-profile</b> <i>dot1p-value</i> ]}<br><b>no dot1p-outer</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>qos>sap-egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command explicitly defines the egress outer or single VLAN tag IEEE 802.1p (dot1p) bits marking for <i>fc-name</i>. When the marking is set, all packets of <i>fc-name</i> 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 <i>dot1p-value</i>. If the egress packets for <i>fc-name</i> are not IEEE 802.1q or IEEE 802.1p encapsulated, this command has no effect.</p> <p>The optional <b>in-profile</b> <i>dot1p-value</i> <b>out-profile</b> <i>dot1p-value</i> [<b>exceed-profile</b> <i>dot1p-value</i>] parameters on the dot1p-outer command add the capability to mark the in, out, and exceed-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, out, and exceed-profile status needs to be communicated to an access network or customer device that does not support the DE bit.</p> <p>When the <b>in-profile</b> keyword is added, the <b>out-profile</b> keyword must be specified; however, <b>exceed-profile</b> is optional. If the optional <b>exceed-profile</b> <i>dot1p-value</i> is not included, any exceed-profile traffic will be marked with the same dot1p value as configured for the out-of-profile traffic. All inplus-profile traffic is marked with the same value as in-profile traffic.</p> <p>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, out, or exceed-profile status may be indicated via the setting of the DE bit setting if the <b>de-mark(-outer)</b> command is used. The DE value used for inplus is the same as that used for in-profile and the one used for exceed-profile is the same as that used for out of profile.</p> <p>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.</p> <p>The two versions of the command (with and without parameters) are mutually exclusive.</p> <p>This command takes precedence over the <b>dot1p</b> command if both are specified in the same policy, and over the default action where the marking is taken from a packet received at ingress.</p> <p>The <b>no</b> form of the command sets the IEEE 802.1p or IEEE 802.1q priority bits to 0.</p> |
| <b>Default</b>     | no dot1p-outer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>dot1p-value</i> — Specifies the 802.1p value to set for in-profile frames in this forwarding class.</p> <p><b>Values</b>      0 to 7</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**in-profile** *dot1p-value* — Specifies the 802.1p value to set for in-profile frames in this forwarding class.

**Values** 0 to 7

**out-profile** *dot1p-value* — Specifies the 802.1p value to set for out-of-profile frames in this forwarding class.

**Values** 0 to 7

**exceed-profile** *dot1p-value* — Specifies the 802.1p value to set for exceed-profile frames in this forwarding class.

**Values** 0 to 7

## dscp

**Syntax** **dscp** {*dscp-name* | **in-profile** *dscp-name* **out-profile** *dscp-name* [**exceed-profile** *dscp-name*]}

**no dscp**

**Context** config>qos>sap-egress>fc

**Description** This command configures a DSCP to be used for remarking packets from the specified FC. If the optional **exceed-profile**, **in-profile**, or **out-profile** keyword is specified, the command will remark different DSCP depending on whether the packet was classified to be exceed, in-profile, or out-of-profile ingress to the node. All inplus-profile traffic is marked with the same value as in-profile traffic.

**Default** no dscp

**Parameters** *dscp-name* — Specifies a DSCP name that has been previously mapped to a value using the **dscp-name** command. The DSCP 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

**exceed-profile** *dscp-name* — This optional parameter specifies the DSCP name to be used to remark the traffic that is exceed-profile. If not specified, this defaults to the same value configured for **out-profile** parameter.

**in-profile** *dscp-name* — Specifies the DSCP name to be used to remark the traffic that is in-profile.

**out-profile** *dscp-name* — Specifies the DSCP name to be used to remark the traffic that is out-of-profile.

---

## hsmda

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hsmda</b>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>qos>sap-egress>fc                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | 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 HSMDA queue mappings for the same forwarding class; 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

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>queue <i>hsmda-queue-id</i> [port-redirect-group-queue]</b><br><b>no queue</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>qos>sap-egress>fc>hsmda                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | 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 <b>queue <i>queue-id</i></b> 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 listed in [Table 46](#). 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 46** 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                |



**Table 46 Default FC HSM DA Queue ID Mappings (Continued)**

| Forwarding Class | Default HSM DA Queue ID |
|------------------|-------------------------|
| L2               | queue 2                 |
| BE               | queue 1                 |

Table 47 describes how packets are mapped to queues based on the type of service and the various forwarding types.

**Table 47 Ingress HSM DA 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                                       |
| 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 FDB |
| VRPN                                    | All packets matching the FC | Packets with Broadcast DA | IP Multicast Packets      | None                                       |

The forwarding class queue mappings may be modified at any time. The sub-forwarding classes inherit the parent forwarding classes queue mappings.

The **no** form of this command returns the HSM DA queue mapping for the queue to the default mapping for the forwarding class.

**Parameters** *hsm da-queue-id* — Configures a specific HSM DA queue.

**Values**

1 to 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

**port-redirect-group-queue** — 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 known as SAP-based redirection.

policer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>policer</b> <i>policer-id</i> [[[ <b>port-redirect-group-queue</b> ] [ <b>queue</b> <i>queue-id</i> ]   <b>group</b> <i>group-name</i> [ <b>instance</b> <i>instance-id</i> ] [ <b>queue</b> <i>queue-id</i> ]]]<br><b>no policer</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-egress>fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | 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 that 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 <i>queue-id</i> , or port-redirect-group queue <i>queue-id</i> , or group <i>queue-group-name</i> instance <i>instance-id</i> queue <i>queue-id</i> commands) or an egress policer ( <b>policer</b> <i>policer-id</i> ). 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 <i>policer-id</i> 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 supports 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 specified in the QoS policy, or the fact that a forwarding class must be redirected is identified in the QoS policy and the specific queue group instance is only identified at the time the QoS policy 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 the 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** *group-name* is not defined as an egress queue-group-template, the **policer** command will fail. Also, 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** *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.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no policer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b> | <p><i>policer-id</i> — When the forwarding class <b>policer</b> command is executed, a valid <i>policer-id</i> must be specified. The parameter <i>policer-id</i> references a <i>policer-id</i> that has already been created within the <b>sap-egress</b> QoS policy.</p> <p><b>Values</b> 1 to 63</p> <p><b>port-redirect-group-queue</b> — 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 the QoS policy is applied to the SAP. Therefore, this parameter is only valid if SAP-based redirection is required.</p> <p><b>queue queue-id</b> — This parameter overrides the forwarding class default egress queue destination to a specified <i>queue-id</i>. If port-redirect-group is not configured, this will be a local SAP queue of that <i>queue-id</i>. A queue of ID <i>queue-id</i> must exist within the egress QoS policy. If <b>port-redirect-group-queue</b> is configured, the <b>queue queue-id</b> in the egress port queue group instance is used.</p> <p><b>Values</b> 1 to 8</p> <p><b>Default</b> Derived from forwarding class assignment in queue-group definition.</p> <p><b>group group-name</b> — The <b>group queue-group-name</b> 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 are applied or the policer command will fail.</p> <p><b>Values</b> Any qualifying egress queue group name</p> <p><b>Default</b> policer-output-queues</p> <p><b>queue queue-id</b> — The <b>queue group-queue-id</b> 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.</p> <p><b>Values</b> 1 to 8</p> <p><b>Default</b> Derived from forwarding class assignment in queue-group definition</p> |

**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 to 40960

**Default** 1

## prec

**Syntax** **prec** {*ip-prec-value* | **in-profile** *ip-prec-value* **out-profile** *ip-prec-value* [**exceed-profile** *ip-prec-value*]}  
**no prec**

**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/exceed-profile is specified, the command will remark different IP precedence values depending on whether the packet was classified to be in, exceed, or out-of-profile. All inplus-profile traffic is marked with the same value as in-profile traffic.

**Parameters** *ip-prec-value* — This parameter specifies the IP precedence to be used to remark all traffic

**Values** 0 to 7

**exceed-profile** *ip-prec-value* — This optional parameter specifies the IP precedence to be used to remark traffic that is exceed-profile. If not specified, this defaults to the same value configured for the **out-profile** parameter.

**Values** 0 to 7

**in-profile** *ip-prec-value* — This parameter specifies the IP precedence to be used to remark traffic that is in-profile.

**Values** 0 to 7

**out-profile** *ip-prec-value* — This parameter specifies the IP precedence to be used to remark traffic that is out-of-profile.

**Values** 0 to 7

## queue

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

**Context** config>qos>sap-egress>fc

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command overrides the default queue mapping for <b>fc</b> <i>fc-name</i>. The specified <i>queue-id</i> must exist within the policy before the mapping can be made. When the forwarding class mapping is executed, all traffic is classified to <i>fc-name</i> on a SAP using this policy.</p> <p>The <b>no</b> form of this command sets the <i>queue-id</i> back to the default queue for the forwarding class (queue 1).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>     | no queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>queue-id</i> — Specifies the SAP egress <i>queue-id</i> to be associated with the forwarding class. The <i>queue-id</i> must be an existing queue defined in <i>sap-egress policy-id</i>.</p> <p><b>Values</b> 1 to 8</p> <p><b>Default</b> 1</p> <p><i>queue-group-name</i> — This optional parameter is used to redirect the forwarding type within the forwarding class to the specified <i>queue-id</i> within the <i>queue-group-name</i>. 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 <i>queue-group-name</i> are configured in the <b>config&gt;qos&gt;queue-group-templates</b> 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.</p> <p><i>instance-id</i> — This parameter is used to specify the specific instance of a queue group with template <i>queue-group-name</i> to which this queue should be redirected. This parameter is only valid for queue groups on egress ports where policy-based redirection is required.</p> <p><b>Values</b> 1 to 40960</p> <p><b>Default</b> 1</p> <p><b>port-redirect-group-queue</b> — This keyword is used to mark a given forwarding class queue for redirection to an egress queue group <i>queue</i>. 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.</p> |

## hsmda-queues

|                    |                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hsmda-queues</b>                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>qos>sap-egress                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command enters 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 unaware of the type of hardware that the SAP or subscriber is traversing.</p> |

## low-burst-max-class

|                    |                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>low-burst-max-class</b> <i>class</i><br><b>no low-burst-max-class</b>                                                                                                               |
| <b>Context</b>     | config>qos>sap-egress>hsmda-queues                                                                                                                                                     |
| <b>Description</b> | This command assigns the low burst maximum class to associate with the HSMDA queue.<br><br>The <b>no</b> form of this command returns the class id for the queue to the default value. |
| <b>Parameters</b>  | <i>class</i> — Specifies the class identifier of the low burst max class for the HSMDA queue.<br><br><b>Values</b> 1 to 32                                                             |

## packet-byte-offset

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>packet-byte-offset</b> { <b>add</b> <i>add-bytes</i>   <b>subtract</b> <i>sub-bytes</i> }<br><b>no packet-byte-offset</b>                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>qos>sap-egress>hsmda-queues                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | 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 14-byte Ethernet DLC header, 4-byte or 8-byte VLAN tag (optional), 20-byte IP header, IP payload, and the 4-byte CRC (everything except the preamble and inter-frame gap). For 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 that 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.

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 HSMDB 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 subscriber packets on an Ethernet aggregation network.

The packet-byte-offset value may be overridden for the HSMDB queue at the SAP or subscriber profile level.

The **no** form of this 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.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><i>add-bytes</i> — Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the <b>add</b> or <b>subtract</b> keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The <b>add</b> keyword is mutually exclusive with the <b>subtract</b> keyword.</p> <p><b>Values</b> 1 to 31</p> <p><i>sub-bytes</i> — Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The <b>subtract</b> keyword is mutually exclusive with the <b>add</b> keyword. Either the <b>add</b> or <b>subtract</b> keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The minimum resulting packet size used by the system is 64 bytes with an HSMDB.</p> <p><b>Values</b> 1 to 64</p> |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

queue

|                |                                    |
|----------------|------------------------------------|
| <b>Syntax</b>  | [no] queue <i>queue-id</i>         |
| <b>Context</b> | config>qos>sap-egress>hsmdb-queues |



**Description** This command, within the QoS policy HSMDBA-queues context, is a container for the configuration parameters controlling the behavior of an HSMDBA queue. Unlike the standard QoS policy queue command, this command is not used to actually create or dynamically assign the queue to the object to 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 HSMDBA SAPs and subscribers, all eight queues exist at the moment the system allocates an HSMDBA 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. HSMDBA queues are handled differently. The create time auto-expedite, explicit expedite, and best-effort qualifiers have been eliminated; 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 HSMDBA Queues**

Another difference between HSMDBA queues and standard service queues is the lack of multipoint queues. At ingress, an HSMDBA SAP or subscriber does not require multipoint 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 hsmdbA-queues node supports a maximum of eight queues.

#### **Every HSMDBA Queue Supports Profile Mode Implicitly**

Unlike standard service queues, the HSMDBA 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 dependent 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 HSMDBA is possible using the port-redirect-group parameter. If this is specified, packets are redirected to the queue-id in the HSMDBA queue group instance named at the time the egress QoS policy is applied to the SAP.

The **no** form of this command restores the defined queue-id to its default parameters. All HSMMDA queues having the queue-id and associated with the QoS policy are re-initialized to default parameters.

**Parameters** *queue-id* — Defines the context from which the eight ingress or egress queues will be entered for editing purposes.

## adaptation-rule

**Syntax** **adaptation-rule** [**pir** {**max** | **min** | **closest**}]  
**no adaptation-rule**

**Context** config>qos>sap-egress>hsmmda-queue>queue

**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 **no** form of this command is used to return the policer's metering and profiling hardware adaptation rules to closest.

**Default** closest

**Parameters** **max** — Specifies 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** — Specifies 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** — Specifies 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.

## burst-limit

**Syntax** **burst-limit** *size* [**bytes** | **kilobytes**]  
**no burst-limit**

**Context** config>qos>sap-egress>hsmmda-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.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><b>size</b> — 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.</p> <p><b>Values</b> 1 to 1000000</p> <p><b>bytes</b> — Specifies that the value given for size must be interpreted as the burst limit in bytes.</p> <p><b>kilobytes</b> — Specifies that the value given for size must be interpreted as the burst limit in kilobytes.</p> <p><b>Default</b> kilobytes</p> |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## rate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>rate</b> {<i>pir-rate</i>   <b>max</b>}</p> <p><b>no rate</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>qos>sap-egress>hsmda-queues>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command configures a rate limit (PIR) for scheduling packets out of the queue and an optional CIR rate used to determine the profile of packets scheduled from the queue. Configuring a rate limit for a queue on an HSMDA is optional; the default rate is set to maximum (max) causing the shaper to have no effect. The cir keyword is used to configure the rate threshold between the in-profile and out-of-profile state of the queue during scheduling from the queue.</p> |

Since the CIR leaky bucket is updated during scheduling events and not enqueueing events, the profiling function is not based on packet arrival. Instead, the queue absorbs bursts that exceed the queues forwarding rate. In this case, burst tolerance is more heavily affected by the maximum queue depth (mbs) and the PIR shaping behavior than the CIR leaky bucket behavior.

### Egress Profiling-Based Dot1p Remarking

HSMDA egress queues are capable of remarking dot1p and DEI bits based on the current state of the queues CIR. Egress dot1p remarking is enabled at the forwarding class level. Using egress profiling-based dot1p remarking, either two distinct dot1p values may be used to distinguish in-profile and out-of-profile, or just the DEI bit may be toggled.

**SAP and Subscriber Queue Rate Overrides**

The shaping rate and CIR values may be overridden on each SAP or subscriber to which the QoS policy is associated.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><i>pir-rate</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the <b>rate</b> command is executed, a valid PIR setting must be explicitly defined. When the <b>rate</b> command has not been executed, the default PIR of <b>max</b> is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p>The actual PIR rate is dependent on the queue's <b>adaptation-rule</b> parameters and the actual hardware where the queue is provisioned.</p> |
| <b>Values</b>     | 1 to 100000000, <b>max</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>    | <b>max</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**slope-policy**

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>slope-policy</b> <i>hsmda-slope-policy-name</i></p> <p><b>no slope-policy</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>qos>sap-egress>hsmda-queues>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command associates an existing HSMDB slope policy with the QoS policy HSMDB queue. The specified hsmdb-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. When 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.</p> |

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 hsmdda-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 hsmdda-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 this 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. When a slope policy is associated with an HSMDB queue, the policy cannot be deleted.

**Default** hsmda-default

## wrr-weight

**Syntax** **wrr-weight** *weight*  
**no wrr-weight**

**Context** config>qos>sap-egress>hsmdb-queues>queue

**Description** This command assigns the weight value to the HSMDB queue.  
The **no** form of this command returns the weight value for the queue to the default value.

**Parameters** *weight* — Specifies the weight for the HSMDB queue.

**Values** 1 to 32

## wrr-policy

**Syntax** **wrr-policy** *wrr-policy-name*  
**no wrr-policy**

**Context** config>qos>sap-egress>hsmdb-queues

**Description** This command associates an existing HSMDB weighted-round-robin (WRR) scheduling loop policy with the HSMDB queue.

**Parameters** *wrr-policy-name* — Specifies the existing HSMDB WRR policy name to associate with the queue.

## action

**Syntax** **action** [*fc fc-name*] [**hsmdb-counter-override** *counter-id*] [**profile** {*in* | *out* | **exceed** | **inplus**}] [**policer** *policer-id*] [**port-redirect-group-queue**] [*queue queue-id*] [**use-fc-mapped-queue**]  
**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 only 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, or profile or HSMDB or 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 HSMDB 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.

It is also possible to redirect the egress packet to a configured policer. The forwarding class or profile can also be optionally specified, but redirection to a policer is mutually exclusive with the **hsmdb-counter-override** keyword.

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 exit 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 **hsmdb-counter-override** keyword is optional. When specified and the egress SAP is created on an HSMDB, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMDB uses each queue's default queue counters for packets mapped to the queue. The **hsmdb-counter-override** keyword is used to map the packet to an explicit exception counter. One of eight counters can be used. When the packet is mapped to an exception counter, the packet will not increment the queue's discard or forwarding counters; instead, the exception discard and forwarding counters will be used. This keyword is mutually exclusive with the redirection to a policer.

The **policer** keyword is optional. When specified, the egress packet will be redirected to the configured policer. Optional parameters allow the user to control how the forwarded policed traffic exits the egress port. By default, the policed forwarded traffic will use a queue in the egress port's policer-output-queue queue group; alternatively, a queue in an instance of a user-configured queue group can be used or a local SAP egress queue. This keyword is mutually exclusive with the **hsmda-counter-override** keyword.

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.

**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 |                                |
|    | class | be, l2, af, l1, h2, ef, h1, nc |

**profile {in | out | exceed | inplus}** — The profile reclassification action is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to the configured 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 exclusive to the **exceed**, **inplus**, and **out** parameters following the profile reclassification action keyword. **In**, **exceed**, **inplus**, 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 exclusive to the **exceed**, **inplus**, and **in** parameters following the profile reclassification action keyword. **In**, **exceed**, **inplus**, 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.

**exceed** — The **exceed** parameter is mutually exclusive to the **out**, **inplus**, and **in** parameters following the profile reclassification action keyword. **In**, **exceed**, **inplus**, or **out** must be specified when the profile keyword is present. When **exceed** is specified, any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.



**inplus** — The **inplus** parameter is mutually exclusive to the **out**, **exceed**, and **in** parameters following the profile reclassification action keyword. **In**, **exceed**, **inplus**, or **out** must be specified when the profile keyword is present. When **inplus** is specified, any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

**hsmda-counter-override** *counter-id* — The **hsmda-counter-override** parameter is optional and only has significance on SAPs that 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** parameter may be overwritten by an IP criteria reclassification rule match. The *counter-id* must be specified as an integer from one to eight. 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** parameter defined. This keyword is mutually exclusive with the redirection to a policer.

**Values** 1 to 8

**policer** *policer-id* — When the action 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 egress QoS policy.

**Values** 1 to 63

**port-redirect-group-queue** *queue-id* — 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 the QoS policy is applied to the SAP. Therefore, this parameter is only valid if SAP-based redirection is required. The queue parameter overrides the policer's default egress queue destination to a specified queue-id in the egress port queue group instance.

**Values** 1 to 8

**queue** *queue-id* — This parameter overrides the policer's default egress queue destination to a specified local SAP queue of that queue-id. A queue of ID queue-id must exist within the egress QoS policy.

**Values** 1 to 8

**use-fc-mapped-queue** — This parameter overrides the policer's default egress queue destination to the queue mapped by the traffic's forwarding class.

## parent-location

**Syntax** **parent-location** {default | sla}  
**no parent-location**

**Context** config>qos>sap-egress

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>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.</p> <p>If a parent scheduler name does not exist at the specified location, the queue will not be parented and will be orphaned.</p> <p>The <b>no</b> form of this command reverts to the default.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>     | parent-location default                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <p><b>default</b> — 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.</p> <p>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 multiservice site.</p> <p><b>sla</b> — 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.</p> <p>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. This parameter is not supported when <b>policers-hqos-manageable</b> is configured in the SAP egress QoS policy.</p> |

## policer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>policer</b> <i>policer-id</i> [<b>create</b>]</p> <p><b>no policer</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>qos>sap-egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>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.</p> <p>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 supports egress policers. By default, the system uses the forwarding class mappings within this queue group to decide which queue within the</p> |

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 specified in the QoS policy, or the fact that a forwarding class must be redirected is identified in the QoS policy and the specific queue group instance is only identified at the time the QoS policy 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 the 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** *group-name* is not defined as an egress queue-group-template, the **policer** command will fail. Also, 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** *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.

|                   |                                                                                                                                                                                                                                                                      |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no policer                                                                                                                                                                                                                                                           |
| <b>Parameters</b> | <i>policer-id</i> — When the forwarding class <b>policer</b> command is executed, a valid <i>policer-id</i> must be specified. The parameter <i>policer-id</i> references a <i>policer-id</i> that has already been created within the <b>sap-egress</b> QoS policy. |
| <b>Values</b>     | 1 to 63                                                                                                                                                                                                                                                              |

## enable-dscp-prec-remarking

|                    |                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] enable-dscp-prec-remarking</b>                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-egress>policer                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command enables DSCP/precedence remarking based on the profile state of a packet being forwarded by a SAP or subscriber egress policer. The DSCP/precedence can be remarked to a value independent of, or separately based on, the packet's profile, if the packet has an exceed, in-profile, or out-of-profile state. |
| <b>Default</b>     | no enable-dscp-prec-remarking                                                                                                                                                                                                                                                                                               |

## enable-exceed-pir

|                    |                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] enable-exceed-pir</b>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>sap-egress>policer                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command enables the forwarding of packets with an exceed-profile state and traffic exceeding the PIR for a SAP egress or a network egress queue group (configured in the egress queue group template) policer. This traffic is forwarded as exceed-profile instead of being dropped. This parameter is not supported when <b>policers-hqos-manageable</b> is configured in the SAP egress QoS policy. |
| <b>Default</b>     | no enable-exceed-pir                                                                                                                                                                                                                                                                                                                                                                                       |

## port-parent

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>port-parent</b> [ <i>weight weight</i> ] [ <i>level level</i> ] [ <i>cir-weight cir-weight</i> ] [ <i>cir-level cir-level</i> ]<br><b>no port-parent</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-egress>policer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command specifies whether this SAP egress policer feeds off a port-level scheduler. When configured, the policer is parented by a port-level scheduler. This requires that <b>policers-hqos-manageable</b> be configured in the SAP egress QoS policy. This command and the SAP egress policer <b>scheduler-parent</b> and the <b>parent</b> commands are mutually exclusive.</p> <p>The <b>port-parent</b> command defines a child/parent association between an egress policer and a port-based scheduler or between an intermediate service scheduler and a port-based scheduler. The <b>port-parent</b> command allows for a set of within-CIR and above-CIR parameters that define the port priority levels and weights for the policer. If the <b>port-parent</b> command is executed without any parameters, the default parameters are used.</p> <p>In this context, the <b>port-parent</b> command and the <b>scheduler-parent</b> command (used to create a parent/child association between a queue and an intermediate scheduler) are mutually exclusive. Executing a <b>port-parent</b> command when a <b>scheduler-parent</b> definition exists causes the current intermediate scheduler association to be removed and replaced by the defined port-parent association. Executing a <b>scheduler-parent</b> command when a <b>port-parent</b> definition exists causes the port scheduler association to be removed and replaced by the defined intermediate scheduler association.</p> <p>Changing the parent context on a SAP egress policy policer may cause a SAP or subscriber or a multiservice site context of the policer (policy associated with a SAP or subscriber profile or a multiservice site) to enter an orphaned state. If an instance of a policer is created on a port or channel that does not have a port scheduler enabled, and the SAP egress policy creating the policer has a port parent association, the policer 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 policer is on a port or channel that has a port scheduler configured and the SAP egress policy defines the policer as having a non-existent intermediate scheduler parent, the policer will be treated as an orphan and will be handled according to the current orphan behavior on the port scheduler.</p> <p>The <b>no</b> form of this command removes a port scheduler parent association for the policer. When removed, if a port scheduler is defined on the port on which the policer instance exists, the policer will be treated as orphaned to the port scheduler.</p> |
| <b>Default</b>     | no port-parent                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <b>weight weight</b> — Defines the weight that the policer will use at the above-CIR port priority level (defined by the <b>level</b> parameter).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

All weight values from all weighted active policers, queues, and schedulers with a common port parent are added together. Then, each individual active weight is divided by the total to determine the percentage of remaining bandwidth provided to the policer, queue, or scheduler after the higher priority level children have been serviced. A weight is considered to be active when the applicable policer, queue, or scheduler has not reached its maximum rate and still has packets to transmit.

The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the **weight** parameter is set to a value of 0, the policer receives bandwidth only after other children with a non-zero weight at this level.

**Values** 0 to 100

**Default** 1

**level** *level* — Defines the port priority that the policer will use to receive bandwidth for its above-CIR offered load.

**Values** 1 to 8 (8 is the highest priority)

**Default** 1

**cir-weight** *cir-weight* — Defines the weight that the policer will use at the within-CIR port priority level (defined by the **cir-level** parameter).

All **cir-weight** values from all weighted active policers, queues, and schedulers with a common port parent are added together. Then, each individual active weight is divided by the total to determine the percentage of remaining bandwidth provided to the policer, queue, or scheduler after the higher priority level children have been serviced. A weight is considered to be active when the applicable policer, queue, or scheduler has not reached its maximum rate and still has packets to transmit.

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 policer receives bandwidth only after the other children with a non-zero weight at this level.

**Values** 0 to 100

**Default** 0

**cir-level** *cir-level* — Defines the port priority that the policer will use to receive bandwidth for its within-CIR offered load. If the **cir-level** parameter is set to a value of 0 (the default value), the policer does not receive bandwidth during the port schedulers within-CIR pass and the **cir-weight** parameter is ignored. If the **cir-level** parameter is 1 or greater, the **cir-weight** parameter is used.

**Values** 0 to 8 (8 is the highest priority)

**Default** 0

## 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 (profile 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.

## scheduler-parent

**Syntax** **scheduler-parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]  
**no scheduler-parent**

**Context** config>qos>sap-egress>policer

**Description** This command defines an optional parent scheduler that governs the available bandwidth given to a policer in addition to the policer's PIR setting. When multiple schedulers, queues, and/or policers share a child status with the parent scheduler, the **weight** or **level** parameters define how this policer contends with the other children for the parent's bandwidth. This command and the configuration of a SAP egress policer **port-parent** or **parent** arbiter are mutually exclusive.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the policer. Scheduler names are configured in the **config>qos>scheduler-policy>tier** context. Multiple schedulers can exist in different scheduler policies with the same *scheduler-name*; in this command, the associated *scheduler-name* pertains to a scheduler that should exist on the egress SAP as the policy is applied and the policer created. When the policer is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly applied or indirectly applied (through a multiservice customer site) to the egress SAP. If the *scheduler-name* does not exist, the policer is placed in the orphaned operational state. The policer will accept packets but will not be bandwidth-limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The SAP to which the policer belongs displays an orphan policer status with the SapEgressPolicerMismatch flag in the **show service sap-using** output. The orphaned state of the policer is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable by the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the policer enters the orphaned state and automatically returns to normal operation when the parent scheduler is available again.

The **no** form of this command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and no error message is returned. When a parent association has been removed, the former child policer attempts to operate based on its configured rate parameter.

Removing the parent association on the policer within the policy takes effect immediately on all policers using the SAP egress QoS policy.

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no scheduler-parent                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b> | <p><b><i>scheduler-name</i></b> — Scheduler names are configured in the <b>config&gt;qos&gt;scheduler-policy&gt;tier</b> context. There are no checks performed at the time of definition to ensure that the <i>scheduler-name</i> exists within an existing scheduler policy. For the policer to use the defined <i>scheduler-name</i>, the scheduler must exist on each egress SAP that the policer is created on. If a <i>scheduler-name</i> does not exist on the egress SAP, the policer operates in an orphaned state.</p> <p><b>Values</b> Any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><b>Default</b> None. Each parental association must be explicitly defined.</p> <p><b>weight <i>weight</i></b> — Defines the relative weight of this policer in comparison with other child policers, queues, and schedulers when competing for bandwidth on the parent <i>scheduler-name</i> at the above-CIR priority level defined by the <b>level</b> parameter. All weight values from all weighted active policers, queues, and schedulers with a common parent are added together. Then, each individual active weight is divided by the total to determine the percentage of remaining bandwidth provided to the policer, queue, or scheduler after the higher priority level children have been serviced. A weight is considered to be active when the applicable policer, queue, or scheduler has not reached its maximum rate and still has packets to transmit. All child policers, queues, and schedulers with a weight of 0 are considered to have the lowest priority at the configured level and are not serviced until all non-zero weighted policers, queues, and schedulers at that level are operating at the maximum bandwidth or are idle.</p> <p><b>Values</b> 0 to 100</p> <p><b>Default</b> 1</p> <p><b>level <i>level</i></b> — Defines the level of hierarchy when compared with other policers, queues, and schedulers when competing for bandwidth on the parent <i>scheduler-name</i>. Children of the parent scheduler with a lower priority will not receive bandwidth until all children with a higher priority have either reached their maximum bandwidth or are idle. Children with the same level are serviced in relation to their relative weights.</p> <p><b>Values</b> 1 to 8 (8 is the highest priority)</p> <p><b>Default</b> 1</p> <p><b>cir-weight <i>cir-weight</i></b> — Defines the relative weight of this policer in comparison with other child policers, queues, or schedulers competing for bandwidth on the parent <i>scheduler-name</i> at the within-CIR priority level defined by the <b>cir-level</b> parameter. All <b>cir-weight</b> values from all weighted active policers, queues, and schedulers with a common parent are added together. Then, each individual active weight is divided by the total to determine the percentage of remaining bandwidth provided to the policer, queue, or scheduler after the higher priority level children have been serviced. A weight is considered to be active when the applicable policer, queue, or scheduler has not reached its maximum rate and still has packets to transmit.</p> |



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 policer receives bandwidth only after the other children with a non-zero weight at this level.

**Values** 0 to 100

**Default** 1

**cir-level** *cir-level* — Defines the level of hierarchy when compared with other policers, queues, and schedulers that the policer will use to receive bandwidth for its within-CIR offered load. If the **cir-level** parameter is set to a value of 0 (the default value), the policer does not receive bandwidth during the schedulers within-CIR pass and the **cir-weight** parameter is ignored. If the **cir-level** parameter is 1 or greater, the **cir-weight** parameter is used.

**Values** 0 to 8 (8 is the highest priority)

**Default** 0

## stat-mode

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>stat-mode {no-stats   minimal   offered-profile-no-cir   offered-profile-cir   offered-total-cir   offered-limited-capped-cir   offered-profile-capped-cir   offered-total-cir-exceed   offered-four-profile-no-cir   offered-total-cir-four-profile}</b><br><b>no stat mode</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>qos>sap-egress>policer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>The <b>sap-egress</b> QoS policy's policer <b>stat-mode</b> command is used to configure the forwarding plane counters that allow offered, forwarded, and dropped 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, out-of-profile, and exceed-profile due to egress profile overrides) 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, for example, will not be configured with a CIR profiling rate and not all policers will receive explicitly reprofiled offered packets. The <b>stat-mode</b> 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.</p> |

While a **no-stats** mode is supported that prevents any packet accounting, the use of the policer's **parent** command requires that the policer's **stat-mode** 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. When 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. The total, allocated, and free statistics can be viewed by using the **tools dump resource-usage card fp** 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 egress policer stat-modes are described in [Table 48](#).

**Table 48 Egress Policer Stat-mode Summary**

| Stat Mode                  | Stat Resources | Traffic Counters (Packet/Octets)                                                                                                  |                                                      | Comments                                                                                                                                                                                        |
|----------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                            |                | Offered                                                                                                                           | Dropped/<br>Forwarded                                |                                                                                                                                                                                                 |
| no-stats                   | 0              | None                                                                                                                              | None                                                 |                                                                                                                                                                                                 |
| minimal                    | 1              | Single counter entering policer                                                                                                   | Single counter for dropped/forwarded exiting policer |                                                                                                                                                                                                 |
| offered-profile-no-cir     | 2              | In or out entering policer                                                                                                        | In/out entering policer                              | Intended for when the policer does not change the profile of packets. Includes only in-profile and out-of-profile.                                                                              |
| offered-profile-cir        | 4              | In, out, or uncolored (which corresponds to hard in-profile, hard out-of-profile, or soft in- or out-of-profile) entering policer | In/out exiting policer                               | Intended for when the policer can change the profile of packets to in-profile and out-of-profile.                                                                                               |
| offered-total-cir          | 2              | Single counter entering policer                                                                                                   | In/out exiting policer                               |                                                                                                                                                                                                 |
| offered-limited-capped-cir | 4              | In or out entering policer                                                                                                        | In/out exiting policer                               | Intended for when the policer has <b>profile-capped</b> configured. The information is limited compared to <b>offered-profile-capped-cir</b> with the benefit of using one less stat resources. |

**Table 48 Egress Policer Stat-mode Summary (Continued)**

| Stat Mode                      | Stat Resources | Traffic Counters (Packet/Octets)                                                                                                  |                                            | Comments                                                                                                                                                                                |
|--------------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                |                | Offered                                                                                                                           | Dropped/<br>Forwarded                      |                                                                                                                                                                                         |
| offered-profile-capped-cir     | 5              | In, out, or uncolored (which corresponds to hard in-profile, hard out-of-profile, or soft in- or out-of-profile) entering policer | In/out exiting policer                     | Intended for when the policer has <b>profile-capped</b> configured                                                                                                                      |
| offered-total-cir-exceed       | 3              | Single counter entering policer                                                                                                   | In/out/exceed exiting policer              | Intended for when the policer is configured with <b>enable-exceed-pir</b> to forward packets that exceed its configured PIR or when traffic is reclassified at egress to exceed-profile |
| offered-four-profile-no-cir    | 4              | Inplus, in, out, or exceed entering policer                                                                                       | Inplus/in/out/exceed entering policer      | Intended to be used when the policer does not change the profile of the packets and traffic is reclassified at egress to inplus and/or exceed-profile                                   |
| offered-total-cir-four-profile | 4              | Single counter entering policer                                                                                                   | Inplus, in, out, or exceed exiting policer | Intended to be used when the policer can change the profile of the packet and traffic is reclassified at egress to profile inplus                                                       |

When a policer is created within the policy, the default setting for **stat-mode** 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, no statistics are generated.

**minimal** — Counter resource allocation: 1

This stat-mode provides the minimal accounting resource usage and counter information, and includes only the total offered, dropped and forwarded packet and octet counters for traffic entering (offered) and exiting (dropped/forwarded) the policer.

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 and do not count different profile output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate or using exceed PIR.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 49](#).

**Table 49**    **Egress Accounting Statistics Collected in minimal stat-mode**

| Show Output | Accounting Stats Collected |                     |
|-------------|----------------------------|---------------------|
|             | Field                      | Field Description   |
| Off. All    | apo                        | AllPacketsOffered   |
|             | aoo                        | AllOctetsOffered    |
| Dro. All    | apd                        | AllPacketsDropped   |
|             | aod                        | AllOctetsDropped    |
| For. All    | apf                        | AllPacketsForwarded |
|             | aof                        | AllOctetsForwarded  |

**offered-profile-no-cir** — Counter resource allocation: 2

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering the policer. inplus-profile traffic is counted with the in-profile counters and exceed-profile traffic is counted with the out-of-profile counters.

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, dropped, and forwarded stats are required from the egress policer, but a CIR or **enable-exceed-pir** 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 or using **enable-exceed-pir**.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 50](#).

**Table 50 Egress Accounting Statistics Collected in offered-profile-no-cir stat-mode**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. InProf  | ipo                        | InProfilePacketsOffered      |
|              | ioo                        | InProfileOctetsOffered       |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered   |
|              | ooo                        | OutOfProfileOctetsOffered    |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-profile-cir** — Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is made up of traffic that is in-plus-profile, in-profile, out-of-profile, exceed-profile, soft in-profile, and soft out-of-profile. The offered counters cover traffic reclassified to in-profile (which includes traffic reclassified to in-plus-profile), traffic reclassified to out-of-profile (which includes traffic reclassified to exceed-profile) and traffic which has not been reclassified at egress (Uncolor). In the dropped and forwarded counters, in-plus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

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 profile-based offered, dropped and forwarded 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.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 51](#).

**Table 51 Egress Accounting Statistics Collected in offered-profile-cir stat-mode**

| Show Output  | Accounting Stats Collected |                            |
|--------------|----------------------------|----------------------------|
|              | Field                      | Field Description          |
| Off. InProf  | ipo                        | InProfilePacketsOffered    |
|              | ioo                        | InProfileOctetsOffered     |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered |
|              | ooo                        | OutOfProfileOctetsOffered  |
| Off. Uncolor | ucp                        | UncoloredPacketsOffered    |
|              | uco                        | UncoloredOctetsOffered     |
| Dro. InProf  | ipd                        | InProfilePacketsDropped    |
|              | iod                        | InProfileOctetsDropped     |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped |
|              | ood                        | OutOfProfileOctetsDropped  |
| For. InProf  | ipf                        | InProfilePacketsForwarded  |
|              | iof                        | InProfileOctetsForwarded   |

**Table 51 Egress Accounting Statistics Collected in offered-profile-cir stat-mode (Continued)**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-total-cir** — Counter resource allocation: 2

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

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 untrusted 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 untrusted packets.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 52](#).

**Table 52 Egress Accounting Statistics Collected in offered-total-cir stat-mode**

| Show Output  | Accounting Stats Collected |                            |
|--------------|----------------------------|----------------------------|
|              | Field                      | Field Description          |
| Off. All     | apo                        | AllPacketsOffered          |
|              | aoo                        | AllOctetsOffered           |
| Dro. InProf  | ipd                        | InProfilePacketsDropped    |
|              | iod                        | InProfileOctetsDropped     |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped |
|              | ood                        | OutOfProfileOctetsDropped  |

**Table 52**      **Egress Accounting Statistics Collected in offered-total-cir stat-mode (Continued)**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-limited-capped-cir** — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is made up of traffic that is inplus-profile, in-profile, out-of-profile, exceed-profile, soft in-profile, and soft out-of-profile. The offered counters cover in-profile traffic (which includes traffic reclassified to inplus-profile) and out-of-profile traffic (which includes traffic reclassified to exceed-profile). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

**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 **soft in-profile** with **profile in** and **soft-out-of-profile** with **profile out** and eliminates the offered-undefined statistic.

The impact of using **offered-limited-capped-cir** stat-mode while profile-capped mode is disabled is 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.

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 53](#).



**Table 53 Egress Accounting Statistics Collected in offered-limited-capped-cir stat-mode**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. InProf  | ipo                        | InProfilePacketsOffered      |
|              | ioo                        | InProfileOctetsOffered       |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered   |
|              | ooo                        | OutOfProfileOctetsOffered    |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-profile-capped-cir** — Counter resource allocation: 5

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is made up of traffic that is inplus-profile, in-profile, out-of-profile, exceed-profile, soft in-profile, and soft out-of-profile. The offered counters cover traffic reclassified to in-profile (which includes traffic reclassified to inplus-profile), traffic reclassified to out-of-profile (which includes traffic reclassified to exceed-profile) and traffic that has not been reclassified at egress (uncolored). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

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 inplus**, **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 is 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 (hard in-profile) instead of offered-undefined (uncolored).

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 54](#).

**Table 54** Egress Accounting Statistics Collected in offered-profile-capped-cir stat-mode

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. InProf  | ipo                        | InProfilePacketsOffered      |
|              | ioo                        | InProfileOctetsOffered       |
| Off. OutProf | opo                        | OutOfProfilePacketsOffered   |
|              | ooo                        | OutOfProfileOctetsOffered    |
| Off. Uncolor | ucp                        | UncoloredPacketsOffered      |
|              | uco                        | UncoloredOctetsOffered       |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |

**offered-total-cir-exceed** — Counter resource allocation: 3

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter. The **offered-total-cir-exceed** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-total-cir-exceed** mode is similar to the **offered-total-cir** mode except that it includes support for forwarded and dropped counters for **profile exceed**.

This mode is intended to be used when the policer is configured with **enable-exceed-pir** to forward packets that exceed its configured PIR or when traffic is egress reclassified to profile exceed. The mode gives the forwarded and dropped counters per profile (in, out, exceed). It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 55](#).

**Table 55 Egress Accounting Statistics Collected in offered-total-cir-exceed stat-mode**

| Show Output  | Accounting Stats Collected |                              |
|--------------|----------------------------|------------------------------|
|              | Field                      | Field Description            |
| Off. All     | apo                        | AllPacketsOffered            |
|              | aoo                        | AllOctetsOffered             |
| Dro. InProf  | ipd                        | InProfilePacketsDropped      |
|              | iod                        | InProfileOctetsDropped       |
| Dro. OutProf | opd                        | OutOfProfilePacketsDropped   |
|              | ood                        | OutOfProfileOctetsDropped    |
| Dro. ExcProf | xpd                        | ExceedProfilePktsDropped     |
|              | xod                        | ExceedProfileOctetsDropped   |
| For. InProf  | ipf                        | InProfilePacketsForwarded    |
|              | iof                        | InProfileOctetsForwarded     |
| For. OutProf | opf                        | OutOfProfilePacketsForwarded |
|              | oof                        | OutOfProfileOctetsForwarded  |
| For. ExcProf | xpf                        | ExceedProfilePktsForwarded   |
|              | xof                        | ExceedProfileOctetsForwarded |

**offered-four-profile-no-cir** — Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering the policer. Offered, dropped, and forwarded counters are provided for inplus, in, out and exceed-profile traffic.

The **offered-four-profile-no-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-four-profile-no-cir** mode is similar to the **offered-profile-no-cir** mode except that it includes support for offered, dropped, and forwarded counters for both inplus-profile and exceed-profile.

This mode is intended to be used when traffic is egress reclassified to inplus and/or exceed-profile. It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 56](#).

**Table 56** Egress Accounting Statistics Collected in offered-four-profile-no-cir stat-mode

| Show Output     | Accounting Stats Collected |                             |
|-----------------|----------------------------|-----------------------------|
|                 | Field                      | Field Description           |
| Off. InProf     | ipo                        | InProfilePacketsOffered     |
|                 | ioo                        | InProfileOctetsOffered      |
| Off. OutProf    | opo                        | OutOfProfilePacketsOffered  |
|                 | ooo                        | OutOfProfileOctetsOffered   |
| Off. ExcProf    | xpo                        | ExceedProfilePacketsOffered |
|                 | xoo                        | ExceedProfileOctetsOffered  |
| Off. InplusProf | ppo                        | InplusProfilePacketsOffered |
|                 | poo                        | InplusProfileOctetsOffered  |
| Dro. InProf     | ipd                        | InProfilePacketsDropped     |
|                 | iod                        | InProfileOctetsDropped      |
| Dro. OutProf    | opd                        | OutOfProfilePacketsDropped  |
|                 | ood                        | OutOfProfileOctetsDropped   |
| Dro. ExcProf    | xpd                        | ExceedProfilePktsDropped    |
|                 | xod                        | ExceedProfileOctetsDropped  |

**Table 56 Egress Accounting Statistics Collected in offered-four-profile-no-cir stat-mode (Continued)**

| Show Output     | Accounting Stats Collected |                              |
|-----------------|----------------------------|------------------------------|
|                 | Field                      | Field Description            |
| Dro. InprofProf | ppd                        | InplusProfilePktsDropped     |
|                 | pod                        | InplusProfileOctetsDropped   |
| For. InProf     | ipf                        | InProfilePacketsForwarded    |
|                 | iof                        | InProfileOctetsForwarded     |
| For. OutProf    | opf                        | OutOfProfilePacketsForwarded |
|                 | oof                        | OutOfProfileOctetsForwarded  |
| For. ExcProf    | xpf                        | ExceedProfilePktsForwarded   |
|                 | xof                        | ExceedProfileOctetsForwarded |
| For. InplusProf | ppf                        | InplusProfilePktsForwarded   |
|                 | pof                        | InplusProfileOctetsForwarded |

**offered-total-cir-four-profile** — Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. There is a separate dropped and forwarded counter for inplus, in, out and exceed-profile traffic.

The **offered-total-cir-four-profile** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-total-cir-four-profile** mode is similar to the **offered-total-cir** except that it includes support for forwarded and dropped counters for both **profile inplus** and **profile exceed**.

This mode is intended to be used when traffic is reclassified at egress to inplus-profile. It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 57](#).

**Table 57** Egress Accounting Statistics Collected in offered-total-cir-four-profile stat-mode

| Show Output     | Accounting Stats Collected |                              |
|-----------------|----------------------------|------------------------------|
|                 | Field                      | Field Description            |
| Off. All        | apo                        | AllPacketsOffered            |
|                 | aoo                        | AllOctetsOffered             |
| Dro. InProf     | ipd                        | InProfilePacketsDropped      |
|                 | iod                        | InProfileOctetsDropped       |
| Dro. OutProf    | opd                        | OutOfProfilePacketsDropped   |
|                 | ood                        | OutOfProfileOctetsDropped    |
| Dro. ExcProf    | xpd                        | ExceedProfilePktsDropped     |
|                 | xod                        | ExceedProfileOctetsDropped   |
| Dro. InprofProf | ppd                        | InplusProfilePktsDropped     |
|                 | pod                        | InplusProfileOctetsDropped   |
| For. InProf     | ipf                        | InProfilePacketsForwarded    |
|                 | iof                        | InProfileOctetsForwarded     |
| For. OutProf    | opf                        | OutOfProfilePacketsForwarded |
|                 | oof                        | OutOfProfileOctetsForwarded  |
| For. ExcProf    | xpf                        | ExceedProfilePktsForwarded   |
|                 | xof                        | ExceedProfileOctetsForwarded |
| For. InplusProf | ppf                        | InplusProfilePktsForwarded   |
|                 | pof                        | InplusProfileOctetsForwarded |

## policers-hqos-manageable

**Syntax** [no] policers-hqos-manageable

**Context** config>qos>sap-egress

**Description** This command specifies that the policers within this SAP egress policy are to be managed by the Hierarchical QoS (HQoS) process when the policy is applied to either the egress part of a SAP configuration or the egress part of an SLA profile, with multiservice sites (MSS) supported for SAPs. When enabled, egress policers and queues can be managed together in the same HQoS hierarchy.

To be managed by HQoS, egress policers within a SAP egress QoS policy must be configured with either a **scheduler-parent** or **port-parent** command, or be orphaned to an egress port scheduler applied on a Vport or port.

The **policers-hqos-manageable** command and **parent-location sla** or policers with **enable-exceed-pir** or **stat-mode no-stats** within an SAP egress QoS policy are mutually exclusive.

In order to prevent HQoS from measuring the traffic through a policer managed by HQoS, then again through a post-policer access egress queue group queue, post-policer access egress queue groups must be configured with **no queues-hqos-manageable** so that their queues are not managed by HQoS.

A post-policer local queue is not supported with HQoS managed policers, nor are those mapped by the **use-fc-mapped-queue** parameter in a criteria action statement. The **policers-hqos-manageable** command is not supported for SAP egress dynamic policers or on a 7950 XRS.

The **no** form of this command results in policers within this SAP QoS egress policy being non-HQoS-manageable.

**Default** no policers-hqos-manageable

## post-policer-mapping

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>post-policer-mapping</b> <i>mapping-policy-name</i><br><b>no post-policer-mapping</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>qos>sap-egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command applies a post-policer mapping policy in a SAP egress QoS policy. The policy contains forwarding class and profile remapping statements, which remap the forwarding class and profile state of an egress policed packet (the profile being the resulting profile after the packet has been processed by the egress policer) to another forwarding class and profile.</p> <p>The remapping applies to all policers within the SAP egress QoS policy, including regular child policers and policers configured in an IP/IPv6 criteria action statement, except for dynamic policers.</p> <p>Post-policer mapping is supported on FP3- and higher-based hardware, with the exception of the 7750 SR-a4/a8, which does not support egress policers resulting in the policy being ignored.</p> <p>The <b>no</b> form of this command deletes the post-policer mapping policy from the SAP egress QoS policy.</p> |
| <b>Parameters</b>  | <i>mapping-policy-name</i> — Specifies the name of the post-policer mapping policy, up to 32 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

---

## prec

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>prec</b> <i>ip-prec-value</i> [ <b>hsmdda-counter-override</b> <i>counter-id</i> ] [ <b>fc</b> <i>fc-name</i> ] [ <b>profile</b> { <b>in</b>   <b>out</b>   <b>exceed</b>   <b>inplus</b> }]<br><b>no prec</b> <i>ip-prec-value</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>qos>sap-egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>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 HSMDDA 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 HSMDDA queue accounting is to use the counters associated with the queue to which the packet is mapped (applies to the 7450 ESS and 7750 SR).</p> <p>The IP precedence bits used to match against precedence reclassification rules come from the Type of Service (ToS) field within the IPv4 header. If the packet does not have an IPv4 header, precedence-based matching is not performed.</p> <p>The reclassification actions from a precedence reclassification rule may be overridden by a DSCP or IP flow matching event.</p> <p>The <b>fc</b> 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 DSCP, ipv6-criteria, or ip-criteria match occurs after the IP precedence 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 IP precedence match will be used.</p> <p>The <b>profile</b> 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 DSCP, IPv6 criteria, or IP criteria match occurs after the IP precedence 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 IP precedence match will be used.</p> <p>The <b>hsmdda-counter-override</b> keyword applies to the 7450 ESS and 7750 SR and is optional. When specified and the egress SAP is created on an HSMDDA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMDDA uses each queues default queue counters for packets mapped to the queue. The <b>hsmdda-counter-override</b> keyword is used to map the packet to an explicit exception counter. One of eight counters can be used. When the packet is mapped to an exception counter, the packet will not increment the queue's 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 <b>hsmdda-counter-override</b> action defined.</p> <p>The <b>no</b> form of this command removes the reclassification rule from the SAP egress QoS policy.</p> |



- 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, IPv6 criteria, 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 precedence value, the **prec** command must be re-executed without the **fc** parameter defined.
- Values**        be, l1, af, l2, h1, ef, h2 or nc
- profile {in | out | exceed | inplus}** — This keyword is optional. When specified, packets matching the IP precedence value will be explicitly reclassified to the specified profile regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by a higher priority DSCP, IPv6 criteria, or IP criteria reclassification match. To remove the profile reclassification action for the specified precedence value, the **prec** command must be re-executed without the **profile** parameter defined.
- in** — Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.
- out** — Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.
- exceed** — Specifies that any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.
- inplus** — Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.
- hsmdda-counter-override** *counter-id* — The **hsmdda-counter-override** parameter is optional and only has significance on SAPs that 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 **hsmdda-counter-override** parameter may be overwritten by an IP criteria reclassification rule match. The specified *counter-id* must be specified as an integer from one to eight. To remove the HSMDA exception counter reclassification action for the specified DSCP value, the **dscp** command must be re-executed without the **hsmdda-counter-override** parameter defined.
- Values**        1 to 8

## queue

- Syntax**        **queue** *queue-id* [*queue-type*] [**create**]  
                  **no queue** *queue-id*
- Context**        config>qos>sap-egress

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command creates the context to configure an egress service access point (SAP) QoS policy queue.</p> <p>Explicit definition of an egress queue's type is supported. A single egress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or best effort 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 best effort forwarding classes are mapped to the queue (be, af, l1, or l2), the queue is treated as best effort by the hardware schedulers. The queue type must be defined at the time of queue creation within the policy.</p> <p>The <b>no</b> form of this command removes the <i>queue-id</i> from the SAP egress 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.</p>                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><i>queue-id</i> — The ID for the queue, expressed as an integer. The <i>queue-id</i> uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.</p> <p><b>Values</b>      1 to 32</p> <p><i>queue-type</i> — Specifies the method that system uses to service the queue from a hardware perspective. A keyword can be specified at the time the queue is created. If an attempt is made to change the keyword after the queue is initially defined, an error is generated.</p> <p><b>Values</b>      <b>expedite</b> - Specifies that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.</p> <p>                 <b>best-effort</b> - Specifies that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.</p> <p>                 <b>auto-expedite</b> - Allows the system to auto-define the way the queue is serviced by the hardware. When <b>auto-expedite</b> 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.</p> <p><b>Default</b>      auto-expedite</p> <p><b>create</b> — Creates an entry for the queue.</p> |

## adaptation rule

**Syntax**      **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <b>no adaptation-rule</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-egress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>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.</p> <p>When a specific <b>adaptation-rule</b> is removed, the default constraints for <b>pir</b> and <b>cir</b> apply.</p> <p>The <b>no</b> form of this command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>     | adaptation-rule pir closest cir closest                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><b>pir</b> — Defines the constraints enforced when adapting the queue's PIR. The <b>pir</b> parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the <b>pir</b> parameter is not specified, the default constraint applies.</p> <p><b>Values</b></p> <ul style="list-style-type: none"> <li><b>max</b> — Specifies that the operational PIR for the queue will be equal to or less than the requested rate.</li> <li><b>min</b> — Specifies that the operational PIR for the queue will be equal to or greater than the requested rate.</li> <li><b>closest</b> — Specifies that the operational PIR for the queue will be the rate closest to the requested rate.</li> </ul> <p><b>cir</b> — Defines the constraints enforced when adapting the queue's CIR defined. The <b>cir</b> parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the <b>cir</b> parameter is not specified, the default constraint applies.</p> <p><b>Values</b></p> <ul style="list-style-type: none"> <li><b>max</b> — Specifies that the operational rate for the queue will be equal to or less than the requested rate.</li> <li><b>min</b> — Specifies that the operational rate for the queue will be equal to or greater than the requested rate.</li> <li><b>closest</b> — Specifies that the operational rate for the queue will be the rate closest to the requested rate.</li> </ul> |

## avg-frame-overhead

|                |                                                                          |
|----------------|--------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>avg-frame-overhead</b> <i>percent</i><br><b>no avg-frame-overhead</b> |
| <b>Context</b> | 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  $10,000 \times 0.1$  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  $50 \times 20$  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, 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  $500 \times 1.1$  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, 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.

|                   |                                                                                                                                                                                                                                                     |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no avg-frame-overhead                                                                                                                                                                                                                               |
| <b>Parameters</b> | <i>percent</i> — 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. This parameter only applies to the 7450 ESS and 7750 SR. |
| <b>Values</b>     | 0.00 to 100.00                                                                                                                                                                                                                                      |

burst-limit

**Syntax**     **burst-limit {default | size [bytes | kilobytes]}**

**no burst-limit**

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |               |                                               |                |                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|-----------------------------------------------|----------------|---------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>qos>sap-egress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |               |                                               |                |                                                                                             |
| <b>Description</b> | <p>The <b>queue burst-limit</b> command defines 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.</p> <p>The <b>burst-limit</b> command is supported under the <b>sap-ingress</b> and <b>sap-egress</b> QoS policy queues. The command is also supported under the ingress and egress <b>queue-group-templates</b> queues.</p> <p>The <b>no</b> form of this command restores 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 <b>queue-override queue</b> context, any current burst limit override for the queue is removed and the queue's burst limit is controlled by its defining policy or template.</p>                                                                    |               |                                               |                |                                                                                             |
| <b>Default</b>     | no burst-limit                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |               |                                               |                |                                                                                             |
| <b>Parameters</b>  | <p><b>default</b> — Reverts the queue's burst limit to the system default value.</p> <p><b>size</b> — Specifies 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 <b>bytes</b> keyword must be added following <b>size</b>.</p> <table><tr><td><b>Values</b></td><td>1 to 13671 kilobytes<br/>0 to 14,000,000 bytes</td></tr><tr><td><b>Default</b></td><td>No default for <b>size</b>; use the <b>default</b> keyword to specify default burst limit.</td></tr></table> <p><b>bytes</b> — Specifies that the value given for size must be interpreted as the burst limit in bytes.</p> <p><b>kilobytes</b> — Specifies that the value given for size must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is <b>kilobytes</b>.</p> | <b>Values</b> | 1 to 13671 kilobytes<br>0 to 14,000,000 bytes | <b>Default</b> | No default for <b>size</b> ; use the <b>default</b> keyword to specify default burst limit. |
| <b>Values</b>      | 1 to 13671 kilobytes<br>0 to 14,000,000 bytes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |               |                                               |                |                                                                                             |
| <b>Default</b>     | No default for <b>size</b> ; use the <b>default</b> keyword to specify default burst limit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |               |                                               |                |                                                                                             |

**exceed**

|                    |                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>exceed</b>                                                                                                                                                                                                                   |
| <b>Context</b>     | config>qos>sap-egress>queue>drop-tail                                                                                                                                                                                           |
| <b>Description</b> | This command enters the context to configure the queue exceed drop tail parameters. The exceed drop tail defines the queue depth beyond which exceed-profile packets will not be accepted into the queue and will be discarded. |

## high

|                    |                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>high</b>                                                                                                                                                                                                             |
| <b>Context</b>     | config>qos>sap-egress>queue>drop-tail                                                                                                                                                                                   |
| <b>Description</b> | This command enters the context to configure the queue high drop tail parameters. The high drop tail defines the queue depth beyond which in-profile packets will not be accepted into the queue and will be discarded. |

## highplus

|                    |                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>highplus</b>                                                                                                                                                                                                                     |
| <b>Context</b>     | config>qos>sap-egress>queue>drop-tail                                                                                                                                                                                               |
| <b>Description</b> | This command enters the context to configure the queue highplus drop tail parameters. The highplus drop tail defines the queue depth beyond which inplus-profile packets will not be accepted into the queue and will be discarded. |

## percent-reduction-from-mbs

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>percent-reduction-from-mbs</b> <i>percent</i><br><b>no percent-reduction-from-mbs</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>sap-egress>queue>drop-tail>exceed<br>config>qos>sap-egress>queue>drop-tail>high<br>config>qos>sap-egress>queue>drop-tail>highplus<br>config>qos>sap-egress>queue>drop-tail>low                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command configures the egress SAP queue drop tails as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and this percentage is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value and will be discarded.</p> <p>The drop tails apply to packets with the following profile state:</p> <ul style="list-style-type: none"> <li>• Exceed drop tail: exceed-profile</li> <li>• High drop tail: in-profile</li> <li>• Highplus drop tail: inplus-profile</li> <li>• Low drop tail: out-of-profile</li> </ul> |
| <b>Default</b>     | <p>Exceed drop tail: 20%</p> <p>Low drop tail: 10%</p> <p>High drop tail: 0%</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Highplus drop tail: 0%

**Parameters**    *percent* — Specifies the percentage reduction from the MBS for a queue drop tail.  
**Values**        0 to 100, default

# packet-byte-offset

**Syntax**        **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-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 applies for the agg-rate-limit on a SAP, a subscriber, or a multiservice 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 and the packet-byte-offset is not included. However the offsets are applied to the statistics.

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 to 32



*sub-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 from the size of each packet associated with the queue for scheduling and accounting purposes. The minimum resulting packet size used by the system is 1 byte.

**Values** 0 to 64

## percent-rate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>percent-rate</b> <i>pir-percent</i> [ <b>cir</b> <i>cir-percent</i> ] [{ <b>port-limit</b>   <b>local-limit</b> }]<br><b>no percent-rate</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>qos>sap-egress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command configures a queue's PIR and CIR as a percentage of the egress port line rate or as a percentage of its parent scheduler rate or <b>agg-rate</b> rate. When the rates are expressed as a <b>port-limit</b> , 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 Gb and a 10 Gb Ethernet port, the queue's rates will be 10 times greater on the 10 Gb 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 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 or **agg-rate** 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 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. Similarly, 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 any time.

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 (kb/s) if the associated QoS policy was also defined as a rate. Queue-overrides are relative to the limit type specified in the QoS policy.

When configured on an egress HSQ queue group queue, the **cir** keyword is ignored.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the **percent-rate** is performed under the **hs-wrr-group** within the SAP egress QoS policy.

When no **percent-rate** is defined within a SAP egress queue-override, the queue uses the defined shaping rate and CIR within the SAP egress QoS policy associated with the queue.

The **no** form of this command returns the queue to its default shaping rate and CIR.

|                   |                                                                                                                                                                                                                                                                                                                |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>pir-percent</i> — Specifies the queue's PIR as a percentage dependent on the use of the <b>port-limit</b> or <b>local-limit</b> .                                                                                                                                                                           |
|                   | <b>Values</b> 0.01 to 100.00                                                                                                                                                                                                                                                                                   |
|                   | <b>Default</b> 100.00                                                                                                                                                                                                                                                                                          |
|                   | <i>cir-percent</i> — Specifies the queue's CIR as a percentage dependent on the use of the <b>port-limit</b> or <b>local-limit</b> .                                                                                                                                                                           |
|                   | <b>Values</b> 0.00 to 100.00                                                                                                                                                                                                                                                                                   |
|                   | <b>Default</b> 100.00                                                                                                                                                                                                                                                                                          |
|                   | <b>port-limit</b> — Specifies that the configure PIR and CIR percentages are relative to the rate of the port (including the <b>egress-rate</b> setting) to which this queue connects. The <b>port-limit</b> is the default                                                                                    |
|                   | <b>local-limit</b> — Specifies that the configure PIR and CIR percentages are relative to the rate of the queue's parent scheduler <b>rate</b> or <b>agg-rate</b> rate at egress. If there is no parent scheduler rate or <b>agg-rate</b> rate, or those rates are <b>max</b> , the <b>port-limit</b> is used. |

## port-parent

|                    |                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>port-parent</b> [ <i>weight weight</i> ] [ <i>level level</i> ] [ <i>cir-weight cir-weight</i> ] [ <i>cir-level cir-level</i> ]<br><b>no port-parent</b>                                                                                                   |
| <b>Context</b>     | config>qos>sap-egress>queue                                                                                                                                                                                                                                   |
| <b>Description</b> | 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 **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 or multiservice site context of the queue (policy associated with a SAP or subscriber profile or multiservice site) 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 on 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.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no port-parent                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b> | <p><b>weight</b> <i>weight</i> — Defines the weight the queue or scheduler will use at the above-CIR port priority level (defined by the level parameter).</p> <p><b>Values</b> 0 to 100</p> <p><b>Default</b> 1</p> <p><b>level</b> <i>level</i> — Defines the port priority the queue or scheduler will use to receive bandwidth for its above-CIR offered-load.</p> <p><b>Values</b> 1 to 8 (8 is the highest priority)</p> <p><b>Default</b> 1</p> <p><b>cir-weight</b> <i>cir-weight</i> — 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.</p> <p><b>Values</b> 0 to 100</p> <p><b>Default</b> 0</p> |

**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 to 8 (8 is the highest priority)

**Default**      0

rate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rate</b> <i>pir-rate</i> [ <i>cir</i> <i>cir-rate</i> ]<br><b>no rate</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>qos>sap-egress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>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.</p> <p>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.</p> <p>The CIR can be used by the queue's parent commands <i>cir-level</i> and <i>cir-weight</i> parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.</p> <p>The <b>rate</b> command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the <i>queue-id</i>.</p> <p>When configured on an egress HSQ queue group queue, the <b>cir</b> keyword is ignored.</p> <p>This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the rate is performed under the <b>hs-wrr-group</b> within the SAP egress QoS policy.</p> <p>The <b>no</b> form of this command returns all queues created with the <i>queue-id</i> by association with the QoS policy to the default PIR and CIR parameters (<b>max</b>, 0).</p> |
| <b>Default</b>     | rate max cir 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><i>pir-rate</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the <b>rate</b> command is executed, a valid PIR setting must be explicitly defined. When the <b>rate</b> command has not been executed, the default PIR of <b>max</b> is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p>The actual PIR rate is dependent on the queue's <b>adaptation-rule</b> parameters and the actual hardware where the queue is provisioned.</p> <p><b>Values</b> [1 to 200000000   <b>max</b>] kb/s</p> <p><b>Default</b> max</p> <p><i>cir-rate</i> — The <b>cir</b> parameter overrides the default administrative CIR used by the queue. When the <b>rate</b> command is executed, a CIR setting is optional. When the <b>rate</b> command has not been executed or the <b>cir</b> parameter is not explicitly specified, the default CIR (0) is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p><b>Values</b> [0 to 200000000   <b>max</b>] kb/s</p> <p><b>Default</b> 0</p> |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## wred-queue

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>wred-queue</b> [ <i>policy slope-policy-name</i> ] [ <i>mode mode</i> ] [ <i>slope-usage slope-usage</i> ]<br><b>no wred-queue</b>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>qos>sap-egress>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command allows the configuration of WRED per queue with the following options:</p> <ul style="list-style-type: none"> <li>• Native hardware WRED<br/>This uses the hardware per queue WRED capabilities of FP3- and higher-based hardware and is configured with the <b>native</b> keyword.</li> <li>• Pool per queue WRED<br/>This implements each queue in its own pool and uses the WRED capabilities of the pool to provide WRED per queue. This is configured with the <b>pool-per-queue</b> keyword.</li> </ul> |

### Native Hardware WRED

When the **wred-queue mode native** command is configured, the queue uses the WRED capabilities of FP3- and higher-based hardware. In this case, the out-of-profile and exceed-profile traffic map to the low and exceed WRED slopes specified within the slope policy, and the inplus-profile and in-profile traffic uses the MBS drop tail; this requires the **slope-usage** to be configured as **exceed-low**. The instantaneous queue depth is compared against the low and exceed slopes so the time average factor in the slope policy is ignored.

When a policy is not explicitly defined, the default slope policy is used.

When **native** mode is enabled for a queue, the **pool** and **drop-tail** commands are ignored; traffic mapped to a slope that is shut down will use the MBS drop tail.

This is only supported on FP3 hardware.

The **no** form of this command restores the queue default congestion control behavior to the queue.

### Pool-per-queue WRED

When the **wred-queue mode pool-per-queue** command is defined and the queue ID is created, 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 are not changed.

In the case where this is applied with 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 in the **wred-queue-control** context, the queue is 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 forwarding plane. 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 highplus, high, low, and exceed 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 the highplus, high, low, or exceed WRED slope based on the packet's profile. If the packet is inplus-profile, the highplus slope is used. If the packet is in-profile, the high slope is used. If the packet is out-of-profile, the low slope is used. If the packet is exceed-profile, the exceed slope is used. This mapping of packet profile to slope is enabled using the **slope-usage default** parameter. 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, the queue **pool** and **drop-tail** commands are ignored; traffic mapped to a slope that is shut down will use the MBS drop tail.

The resource usage for the WRED queue pool-per-queue per forwarding plane can be seen in the **tools dump resource-usage card [slot-num] fp [fp-number]** output under *Dynamic Q2 Wred Pools*.

The **no** form of this 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 the pool is defined and it exists. The queue then uses the default congestion control behavior.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no wred-queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b> | <p><i>slope-policy-name</i> — Specifies an existing slope policy that is used to override the default WRED slope policy.</p> <p><i>mode</i> — Specifies whether the WRED per queue is using the native FP3- and higher-based hardware WRED capabilities or pool per queue.</p> <p><b>Values</b></p> <p><b>native</b> — uses the hardware per queue WRED capabilities of the FP3- and higher-based hardware and requires <b>slope-usage exceed-low</b>.</p> <p><b>pool-per-queue</b> — each queue uses its own pool and the WRED capabilities of the pool to provide WRED per queue. This is supported on both FP2- and higher-based hardware and requires <b>slope-usage default</b>.</p> <p><b>Default</b> native</p> <p><i>slope-usage</i> — Specifies congestion control to be used.</p> <p><b>Values</b></p> <p><b>default</b> — maps the inplus, in, out, and exceed-profile traffic to the highplus, high, low, and exceed WRED slopes, respectively; this is only supported for <b>pool-per-queue</b> mode.</p> <p><b>exceed-low</b> — maps the out and exceed-profile traffic to the low and exceed WRED slopes with the inplus and in-profile traffic using the MBS drop tail. This is only supported for <b>native</b> mode.</p> <p><b>Default</b> exceed-low</p> |

## scope

|                    |                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>scope {exclusive   template}</b></p> <p><b>no scope</b></p>                                                                                                                                                              |
| <b>Context</b>     | config>qos>sap-egress                                                                                                                                                                                                          |
| <b>Description</b> | <p>Enter the scope of this policy. The scope of the policy cannot be changed if the policy is applied to one or more services.</p> <p>The no form of this command sets the scope of the policy to the default of template.</p> |
| <b>Default</b>     | scope 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.

## 6.6.2.4 Operational Commands

### copy

- Syntax**
- ```
copy hsmdda-pool-policy src-name dst-name [overwrite]
copy hsmdda-scheduler-policy src-name dst-name [overwrite]
copy hsmdda-slope-policy src-name dst-name [overwrite]
copy named-pool-policy src-name dst-name [overwrite]
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**
- hsmdda-pool-policy src-name dst-name** — Indicates that the source HSMDA pool policy ID and the destination policy ID are HSMDA pool 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. This parameter applies only to the 7450 ESS and 7750 SR.
 - hsmdda-scheduler-policy src-name dst-name** — Indicates that the source HSMDA scheduler policy ID and the destination policy ID are HSMDA scheduler 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. This parameter applies only to the 7450 ESS and 7750 SR.
 - hsmdda-slope-policy src-name dst-name** — Indicates that the source HSMDA slope policy ID and the destination policy ID are HSMDA 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. This parameter applies only to the 7450 ESS and 7750 SR.

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
```

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 to 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 to 65535

6.6.2.5 Show Commands



Note: For consistency across 7950 XRS platforms, C-XMAs/XMAs are modeled in SR OS (CLI and SNMP) as MDAs.

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

bcg

Syntax	bcg <i>burst-control-group-name</i> member-queues [at-risk-only] [exp-util-bw <i>megabits-persecond</i>] [port <i>port-id</i>] [sap <i>sap-id</i>] [subscriber <i>sub-ident</i>] [access-queue-grp <i>queue-group-name</i>] [service-id <i>service-id</i>] bcg <i>burst-control-group-name</i> [member-queues] [exp-util-bw <i>megabits-persecond</i>] bcg <i>card-slot-number</i> [ingress egress]
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 1 Mb/s and 1 ms 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 1 Mb/s and the threshold (burst limit) is set to 10 kbytes, the maximum amount of time that can expire before the queue is updated without resulting in a negative bucket depth is 81.92 ms. This can be calculated by taking the number of bits represented by the bucket depth (10 kbytes = $10 * 1,024 * 8 = 81,920$ bits) and dividing it by the rate (81,920 bits / 1,000,000 bits per second = 81.92 ms). 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 10 kbytes deep. If the system visits the queue PIR bucket within 81.92 ms, the resulting decrement operation will leave the bucket. If the system takes longer than 81.92 ms, the decrement result will be greater than 10 kbytes and part of the decrement result will be lost. The net result is that, because of less than timely updates, 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 FP2- and FP3-based Q-chip maintains eight BCGs, each targeting a certain queue bucket visitation time. A 1 s, 40 ms, 20 ms, 10 ms, 5 ms, 1 ms, 500 μ s, and 100 μ s BCG is supported. An FP4-based Q-chip has an additional four BCGs with target visitation times of 50 μ s, 10 μ s, 5 μ s, and 1 μ s. By default, queues are placed in 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 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 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 many Q-chip resources were diverted from BCG visitation processing. This calculated scheduling bandwidth is used to evaluate the worst case 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. This is not applicable to an FP4-based Q-chip, so the **exp-util-bw** parameter is not applicable to an FP4-based Q-chip.

Parameters *burst-control-group-name* — Specifies which globally unique Burst Control Group will be displayed. If the specified Burst Control Group does not exist, the show command will fail.

Values *<slot>/<fp>-<target-visitation-time>[-<direction>]*
slot - 1 to 20
fp - 1 to 8
target-visitation-time -1 μ s, 5 μ s, 10 μ s, 50 μ s, 100 μ s, 500 μ s, 1ms, 5ms, 10ms, 20ms, 40ms, 1s
direction - i, e

member-queues — Includes a list of all queues attached to the specified burst-control-group-name. 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. When the **member-queues** parameter is specified, the output includes a single entry summary of the internal queues.

exp-util-bw *megabits-persecond* — Displays a calculated worst-case visitation rate for the specified BCG name based on the specified value of *megabits-persecond* which also modifies the member queue's 'at-risk' state output. This parameter is not applicable to FP4-based hardware which has dedicated internal resources for handling BCG visitations and is unaffected by scheduling bandwidth.

Values 0 to 4294967

port *port-id* — Specifies that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues associated with the specified *port-id*. If the specified *port-id* does not exist, the show command will fail.

Values Up to 17 characters

sap *sap-id* — Specifies that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues associated with the specified *sap-id*. If the specified *sap-id* does not exist, the show command will fail.

Values Up to 64 characters

subscriber *sub-ident* — Specifies that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues associated with the specified *sub-ident*. If the specified *sub-ident* does not exist, the show command will fail.

Values Up to 32 characters

access-queue-grp *queue-group-name* — Specifies 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*. If the specified *queue-group-name* is not found, the show command will fail.

Values Up to 32 characters

service *service-id* — Specifies that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues associated with the specified *service-id*. If the specified *service-id* does not exist, the show command will fail.

Values 1 to 2147483647

card-slot-number — Specifies that the system should find and display the BCG and calculate the slowest accurate visitation time for all queues associated with the specified *card-slot-number*. If the specified *card-slot-number* does not exist, the show command will fail.

Values 1 to 40

[ingress | egress] — The **ingress** and **egress** keywords specify the direction to display and are mutually exclusive.

queue

Syntax	queue queue-group <i>port-id queue-group-name</i> { ingress egress } [id <i>queue-id</i>] [instance <i>instance-id</i>] queue sap <i>sap-id</i> { ingress egress } [id <i>queue-id</i>] queue subscriber <i>sub-ident-string</i> [sla-profile <i>sla-profile-name</i>] { ingress egress } [id <i>queue-id</i>]
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 displays 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	<i>port-id queue-group-name</i> — Specifies that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues within the specified <i>queue-group-name</i> on the specified <i>port-id</i>. If the specified <i>port-id</i> is not provisioned on the system or the specified <i>queue-group-name</i> is not found on the port's specified direction, the show command will fail. Values port-id: up to 17 characters queue-group-name: up to 32 characters {ingress egress} — Specifies the direction to display and are mutually exclusive. <i>queue-id</i> — Limits the return queues to a single <i>queue-id</i>. Values 1 to 16 <i>instance-id</i> — Specifies which 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 to 65535 <i>sap-id</i> — Specifies that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues within the specified <i>sap-id</i>. If the specified <i>sap-id</i> is not found, the show command will fail. Values up to 64 characters <i>sub-ident-string</i> — Specifies that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues associated with the specified <i>sub-ident-string</i>. If the specified <i>sub-ident-string</i> does not exist, the show command will fail. Values up to 32 characters <i>sla-profile-name</i> — Specifies that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues associated with the specified subscriber and sla-profile <i>sla-profile-name</i>. If the specified <i>sla-profile-name</i> does not exist, the show command will fail.

sap-egress

Syntax	sap-egress [<i>policy-id</i>] [association match-criteria hsmda detail] sap-egress summary
Context	show>qos
Description	This command displays SAP egress QoS policy information.
Parameters	<i>policy-id</i> — Displays information about the specific policy ID. Values 1 to 65535 detail — Displays detailed policy information including policy associations. summary — Displays SAP egress policy summary.

Output The following output is an example of QoS SAP egress information, and [Table 58](#) describes the output fields.

Sample Output

```
*A:PE# show qos sap-egress
```

```
=====
Sap Egress Policies
=====
Policy-Id  Scope      Name          Description
-----
1          Template  default      Default SAP egress QoS policy.
-----
Number of Policies : 1
-----
```

```
*A:PE# show qos sap-egress 1 detail
```

```
=====
QoS Sap Egress
=====
```

```
-----
Sap Egress Policy (1)
-----
```

```
Policy-id           : 1           Scope           : Template
Ethernet-ctag       : False        Parent-loc       : default
Name                : default
Description          : Default SAP egress QoS policy.
Policy Active       : True          Plcrs HQoS Managed : False
Post Plcr Mapping Policy: (Not Specified)
HS Attachment Policy : default
```

```
-----
Dynamic Configuration Information
-----
```

```
PccRule Insert Point : n/a          DynPlcr Insert Point : n/a
CBS                   : Def          MBS                   : Def
Parent                : (Not Specified)
Level                 : 1            Weight                  : 1
Packet Byte Offset    : 0
Stat Mode             : minimal
```

```
-----
Queue Information
-----
```

```
Queue-Id            : 1           Queue-Type          : auto-expedite
Admin PIR            : max         Admin CIR           : 0
PIR Percent          : def         CIR Percent         : def
PIR Rule             : closest     CIR Rule            : closest
CBS                  : Def         MBS                 : Def
Avg Frame Overhead   : 0.00        Packet Byte Offset   : add 0
Named Buffer Pool     : (Not Specified)
Adv Config Policy     : (Not Specified)
Parent               : (Not Specified)
PIR Level/Weight     : 1/1         CIR Level/Weight    : 0/1
```

```

Burst Limit      : default
Wred Queue Mode  : none           Wred Queue SlopeUsage: not-applicable
Slope Policy     : default
High-Plus Drop Tail : def         High Drop Tail      : def
Low Drop Tail    : def           Exceed Drop Tail   : def
  
```

HS Queue Information

```

Queue-Id        : 1
Admin PIR       : max
PIR Percent     : def
PIR Rule        : closest
MBS             : Def
Packet Byte Offset : add 0
Burst Limit     : default
HS Class Weight : 1              HS Wrr Weight      : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool : False
  
```

HS Wrr Group Information

```

HS Wrr Group Id   : 1           Class Weight      : 1
Rate              : max         Percent Rate       : 100.00
PIR Adaptation Rule: closest
HS Wrr Group Id   : 2           Class Weight      : 1
Rate              : max         Percent Rate       : 100.00
PIR Adaptation Rule: closest
  
```

Forwarding Class (FC) Classification and Remarking Information

No Matching Entries

```

FC Queue Queue-Group           InstanceId SapBReDir Plcr
  
```

No Matching Entries

Associations

```

Service-Name      : 1
Service-Id        : 1 (VPLS)      Customer-Id      : 1
- SAP : 2/1/1
- SAP : 5/1/1
  
```

Mirror SAPs

No Mirror SAPs Found.

HSMDA Queue	PIR Admin Rule	Packet Offset	WRR Weight	MBS	Slope Plcy Max Class	WRR Plcy Burst Lmt
1	max	add 0	1	default	default	n/a
	closest				8	default
2	max	add 0	1	default	default	n/a
	closest				8	default
3	max	add 0	1	default	default	n/a
	closest				8	default
4	max	add 0	1	default	default	n/a
	closest				8	default
5	max	add 0	1	default	default	n/a
	closest				8	default
6	max	add 0	1	default	default	n/a
	closest				8	default
7	max	add 0	1	default	default	n/a
	closest				8	default
8	max	add 0	1	default	default	n/a
	closest				8	default

FC HSMDA Queue-id HSMDA Port-Redirect-QGrp-Q

No Matching Entries

DSCP Cntr Id Profile fc

No DSCP-Map Entries Found.

Prec Value Cntr Id Profile fc

No Prec-Map Entries Found.

Dot1p fc Profile

No Dot1p-Map Entries Found.

Match Criteria

No Matching Criteria.

HSMDA Associations

No Associations Found.

=====

*A:PE#

Table 58 Show QoS SAP Egress Output Fields

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Scope	Exclusive — Specifies that this policy can only be applied to a single SAP. Template — Specifies 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. closest — The operational CIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational CIR.
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.

Table 58 Show QoS SAP Egress Output Fields (Continued)

Label	Description (Continued)
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.
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 <i>queue-id</i> that uniquely identifies the queue within the policy.
Explicit/Default	Explicit — Specifies the egress IEEE 802.1p (dot1p) bits marking for <i>fc-name</i>. Default — Specifies that the default dot1p value (0) is used.

Table 58 Show QoS SAP Egress Output Fields (Continued)

Label	Description (Continued)
Service Association	
Service-Id	The unique service ID number that identifies the service in the service domain.
Customer-Id	Specifies the customer ID that 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 that 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.

sap-ingress

Syntax **sap-ingress** [*policy-id*] [**association** | **match-criteria** | **hsmda** | **detail**]
sap-ingress summary

Context show>qos

Description This command displays SAP ingress QoS policy information.

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

Values 1 to 65535

Default all SAP ingress policies

detail — Displays detailed policy information including policy associations.

summary — Displays SAP ingress policy summary.

Output The following output is an example of QoS SAP ingress information, and [Table 59](#) describes the output fields.

Sample Output

```
*A:PE# show qos sap-ingress
```

```
=====
Sap Ingress Policies
=====
Policy-Id  Scope      Name                               Description
-----
```

```

1          Template default          Default SAP ingress QoS policy.
-----
Number of Policies : 1
-----
=====
*A:PE#
*A:PE# show qos sap-ingress 1 detail

=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (1)
-----
Policy-id      : 1                      Scope      : Template
Default FC     : be                    Priority    : Low
Criteria-type  : None
Name           : default
Description    : Default SAP ingress QoS policy.
Policy Active  : True

-----
Dynamic Configuration Information
-----
PccRule Insert Point : n/a          DynPlcr Insert Point : n/a
CBS                  : Def          MBS                  : Def
Parent              : (Not Specified)
Level               : 1              Weight              : 1
Packet Byte Offset  : 0
Stat Mode           : minimal

-----
Q  Mode  CIR Admin  PIR Admin  CBS                      PIR Lvl/Wt Parent
        CIR Rule   PIR Rule   MBS                      CIR Lvl/Wt BurstLimit(B)
        CIR Prcnt  PIR Prcnt  Low Drop Tail
        Named-Buffer Pool      Pkt Bt Ofst      Adv Config Policy Name
        FIR Admin  FIR Rule   Cir-non-profiling
        FIR Prcnt

-----
1  Prio   0          max      def                      1/1      None
        closest    closest    def                      0/1      default
        def        def        def
        (not-assigned)      add 0      (not-assigned)
        0          closest    disabled
        def

11 Prio   0          max      def                      1/1      None
        closest    closest    def                      0/1      default
        def        def        def
        (not-assigned)      add 0      (not-assigned)
        0          closest    disabled
        def

=====
Sap-Ingress FC (Detail)
=====
No FC-Map Entries Found.
=====
-----

```

```

FC    DE-1-out-profile  Profile    In-Remark    Out-Remark    Egr-Override-FC
-----
No SubFC Map Entries Found.

-----
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
-----
Service-Name   : 1
Service-Id     : 1 (VPLS)                Customer-Id   : 1
- SAP : 2/1/1
- SAP : 5/1/1
-----
HSMDA Associations
-----
No Associations Found.

-----
LSP EXP       FC                      Priority
-----
No LspExp-Map Entries Found.

=====
*A:PE#

```

Table 59 Show QoS SAP Ingress Output Fields

Label	Description
Policy-Id	The ID that uniquely identifies the policy.

Table 59 Show QoS SAP Ingress Output Fields (Continued)

Label	Description (Continued)
Scope	Exclusive — Specifies that this policy can only be applied to a single SAP. Template - Specifies 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 enqueueing priority when a packet is marked with a <i>dot1p-value</i> 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).
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 CIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational CIR.
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.

Table 59 Show QoS SAP Ingress Output Fields (Continued)

Label	Description (Continued)
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. 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 enqueueing priority when a packet is marked with a <i>dot1p-value</i> specified.
FC	Specifies the forwarding class overrides.

Table 59 Show QoS SAP Ingress Output Fields (Continued)

Label	Description (Continued)
Priority	The optional priority setting overrides the default enqueueing priority for the packets received on an ingress SAP that uses the policy that matches this rule. High — Specifies that the high enqueueing parameter for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Low — Specifies that the low enqueueing parameter for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested.
DSCP	Specifies the forwarding class or enqueueing 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 <code>fc fc-name</code> parameter is defined on the rule.
Priority	This parameter specifies the default enqueueing priority overrides for all packets received on an ingress SAP using this policy that matches this rule. High — Specifies that the high enqueueing parameter for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Low — Specifies that the low enqueueing parameter for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested.
Prec	Specifies the forwarding class or enqueueing priority when a packet is marked with an IP precedence value (<i>ip-prec-value</i>).
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.

Table 59 Show QoS SAP Ingress Output Fields (Continued)

Label	Description (Continued)
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 enqueueing priority overrides for all packets received on an ingress SAP using this policy.
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 enqueueing priority overrides for all packets received on an ingress SAP using this policy.
Service Association	
Service-Id	The unique service ID number that identifies the service in the service domain.

Table 59 Show QoS SAP Ingress Output Fields (Continued)

Label	Description (Continued)
Customer-Id	Specifies the customer ID that identifies the customer to the service.
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP ingress policy is applied.

6.6.2.6 Debug Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

match-criteria-overload

Syntax	match-criteria-overload match-criteria-overload card <i>slot-number</i> fp <i>fp-number</i> match-criteria-overload card <i>slot-number</i> fp <i>fp-number</i> network <i>network-policy-id</i> direction <i>direction</i> match-criteria-overload card <i>slot-number</i> fp <i>fp-number</i> sap-egress <i>policy-id</i> match-criteria-overload card <i>slot-number</i> fp <i>fp-number</i> sap-ingress <i>policy-id</i>
Context	tools>dump>qos
Description	This command, when used without parameters, allows the operator to identify QoS match-criteria memory banks with one or more policies in overload in the system by slot number, FP number, type (IP-MAC or IPv6), SAP ingress, SAP egress, network, and direction (ingress or egress). This command, when used with the card and fp parameters, lists the QoS policy with match criteria in overload on a specified FP when the optional parameters <i>slot-number</i> and <i>fp-number</i> are specified. This command, when used with the network, sap-egress, or sap-ingress parameters, displays the match-criteria entries not programmed for a specified QoS policy due to an overload condition.
Parameters	card <i>slot-number</i> — Specifies the IOM slot number for the overload command. Values 1 to 10 fp <i>fp-number</i> — Specifies the FP number. Values 1 to 8

network *network-policy-id* — Specifies the network policy ID.

Values 1 to 65538

direction *direction* — Specifies the memory bank ingress or egress direction.

Values ingress or egress

sap-egress *policy-id* — Specifies the SAP policy ID.

Values 1 to 65535

sap-ingress *policy-id* — Specifies the SAP policy ID.

Values 1 to 65535

Output The following output is an example of filter overload status.

Sample Output

```
*A:Dut# tools dump qos match-criteria-overload
=====
QoS Match Criteria Status Overload
=====
Slot  FP  Direction          Type
-----
1     1   Ingress              ip-mac
=====
```

Sample overload output with **card** and **fp** parameters only:

```
*A:Dut# tools dump qos match-criteria-overload card 1 fp 1
=====
QoS Match Criteria Status Overload
=====
Slot  FP  Access/Network  Direction  Type  Policy  # Entries
-----
1     1   Access          Ingress    ip    5        5
=====
```

Sample overload output with **card**, **fp**, {**network** or **sap-egress** or **sap-ingress**} parameters only

```
A:Dut# tools dump qos match-criteria-overload card 1 fp 1 sap-ingress 5
=====
QoS Sap Ingress Match Criteria Status Overload
=====
Slot  FP  Type  Policy          Entry
-----
1     1   IPv4  5                1001
1     1   IPv4  5                1002
=====i
```

7 Queue Sharing and Redirection

7.1 Queue Sharing and Redirection

Queue groups are objects created on the 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 of 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.

7.1.1 Supported Platforms

Queue sharing and redirection is supported on the 7950 XRS, 7750 SR, and 7450 ESS platforms as follows:

- Access SAP port and network port queue groups
- Ingress access and network forwarding plane queue groups except on the 7750 SR-a4/a8

Queue sharing and redirection are also supported in conjunction with the use of a C-XMA, XMA, Ethernet MDA, Ethernet CMA, HSMDA, and VSM MDA.

7.2 Queue Group Applications

7.2.1 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 the 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 be sufficient. 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 preceding 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.

7.2.1.1 Ingress Per SAP Statistics with Ingress Queue Groups

A statistic displaying the number of valid ingress packets received on a SAP, or subscribers on that SAP, is shown in the *sap-stats* output. This is available for SAPs in all services. This is particularly useful to display SAP level traffic statistics when forwarding classes in a SAP ingress policy have been redirected to an ingress queue group.

In the following example, traffic is received on an ingress FP policer with a *packet-byte-offset of subtract 10*. It can be seen that the ingress **queueing stats** and **offered forwarding engine stats** are all zero as the traffic is using the FP ingress policer. The Received Valid statistic is non-zero and matches that seen on the ingress FP queue group, with the difference being that the packet-byte-offset is applied to the **queue group policer octets** but not the **Received Valid** octets.

The value in the Received Valid field may not instantaneously match the sum of the offered stats (even in the case where all traffic is using the SAP queues) when traffic is being forwarded; however, when the traffic has stopped, the Received Valid will equal the sum of the offered stats.

```
*A:PE# show service id 1 sap 1/1/9:1 sap-stats
```

```
=====
Service Access Points(SAP)
=====
Service Id      : 1
SAP             : 1/1/9:1          Encap             : q-tag
Description     : (Not Specified)
Admin State    : Up               Oper State       : Up
Flags          : None
Multi Svc Site : None
Last Status Change : 08/31/2018 11:09:25
Last Mgmt Change  : 08/31/2018 11:09:25
-----
Sap Aggregate Stats
-----
```

```

                                Packets                Octets
Ingress
Aggregate Offered      : 0                0
Aggregate Forwarded    : 0                0
Aggregate Dropped      : 0                0

Egress
Aggregate Forwarded    : 0                0
Aggregate Dropped      : 0                0
-----
Sap Statistics
-----
Last Cleared Time      : 08/31/2018 11:12:17

                                Packets                Octets
CPM Ingress            : 0                0

Forwarding Engine Stats
Dropped                : 0                0
Received Valid         : 5                530
Off. HiPrio            : 0                0
Off. LowPrio           : 0                0
Off. Uncolor           : 0                0
Off. Managed           : 0                0

Queueing Stats(Ingress QoS Policy 10)
Dro. HiPrio            : 0                0
Dro. LowPrio           : 0                0
For. InProf            : 0                0
For. OutProf           : 0                0

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf     : 0                0
Dro. Out/ExcProf       : 0                0
For. In/InplusProf     : 0                0
For. Out/ExcProf       : 0                0
=====
*A:PE#
*A:PE# show card 1 fp 1 ingress queue-group "qg1" instance 1 mode access statistics

=====
Card:1  Acc.QGrp: qg1  Instance: 1
=====
Group Name      : qg1
Description     : (Not Specified)
Pol Ctl Pol    : None                Acct Pol      : None
Collect Stats  : disabled
-----
Statistics
-----
                                Packets                Octets

Ing. Policer: 1  Grp: qg1
(Stats mode: minimal)
Off. All       : 5                530
Dro. All       : 0                0
For. All       : 5                530
=====
*A:PE#

```

7.2.1.2 Ingress Access Port Queue Group Hardware Queue Allocation

When ingress access port queue groups are configured, hardware queues are allocated to each switch fabric destination for each queue configured in the queue group template.

The allocation of ingress access port queue group hardware queues has been optimized for the 7950 XRS-20 systems to avoid allocating ingress hardware queues to XCMs in slots 11 and above.

When the first XCM in slot 11 or above is provisioned, additional ingress hardware queues will be allocated to XCMs in slots 11 to 20 for any configured ingress access port queue group queue. If sufficient hardware queues are unavailable, the XCM provisioning will fail. Adding queues to the queue group template or adding additional ingress access port queue groups will continue to require more hardware queue to be allocated, with the configurations failing if there are not sufficient available. When the last XCM in slot 11 and above is unprovisioned, the related additional hardware queues to all of the XCMs in slots 11 and above are freed.

7.2.2 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.

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.

7.2.3 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/VP RN 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, broadcast, unknown, 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 that 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 that 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.

7.2.4 QoS on Ingress Bindings

Traffic is tunneled between VPRN service instances on different PEs over service tunnels bound to MPLS LSPs or GRE tunnels. The binding of the service tunnels to the underlying transport is achieved either automatically (using the **auto-bind-tunnel** command) or statically (using the **spoke-sdp** command; not on the VPRN IP interface). QoS control can be applied to the service tunnels for traffic ingressing into a VPRN service; see [Figure 18](#).

Figure 18 Ingress QoS Control on VPRN Bindings



An ingress queue group must be configured and applied to the ingress network FP where the traffic is received for the VPRN. All traffic received on that FP for any binding in the VPRN (either automatically or statically configured) that is redirected to a policer in the FP queue group (using **fp-redirect-group** in the network QoS policy) will be controlled by that policer. As a result, the traffic from all such bindings is treated as a single entity (per forwarding class) with regard to ingress QoS control. Any **fp-redirect-group**, **mcast-policer**, **broadcast-policer**, or **unknown-policer** commands in the network QoS policy are ignored for this traffic (IP multicast traffic would use the ingress network queues or queue group related to the network interface).

Ingress classification is based on the configuration of the ingress section of the specified network QoS policy, noting that the dot1p and exp classification is based on the outer Ethernet header and MPLS label whereas the DSCP applies to the outer IP header if the tunnel encapsulation is GRE, or the DSCP in the first IP header in the payload if **ler-use-dscp** is enabled in the ingress section of the referenced network QoS policy.

Ingress bandwidth control does not take into account the outer Ethernet header, the MPLS labels/control word or GRE headers, or the FCS of the incoming frame.

The following command configures the association of the network QoS policy and the FP queue group and instance within the network ingress of a VPRN:

```
configure
  vprn
    network
      ingress
        qos <network-policy-id> fp-redirect-group <queue-group-name>
          instance <instance-id>
```

When this command is configured, it overrides the QoS applied to the related network interfaces for unicast traffic arriving on bindings in that VPRN. The IP and IPv6 criteria statements are not supported in the applied network QoS policy

This is supported for all available transport tunnel types and is independent of the label mode (**vrf** or **next-hop**) used within the VPRN. It is also supported for Carrier-Supporting-Carrier VPRNs.

7.2.5 VXLAN VNI Queue Group Redirection

VXLAN Virtual Network Identifier (VNI) queue group redirection allows IPv4 and IPv6 VXLAN and VXLAN GPE routed transit packets to be mapped to an FP ingress or port egress queue group instance based on the VXLAN VNI. This allows QoS control to be applied to transit VXLAN traffic.

QoS for transit VXLAN can also be achieved by using VNI classification in an **ip-criteria** or **ipv6-criteria** statement within a SAP ingress QoS policy (see [Virtual Network Identifier \(VNI\) Classification](#)). The main difference between VNI SAP ingress classification and queue group redirection is that the former is only for ingress and is less scalable because it uses SAP policers or queues instead of queue groups; however it does support ranges of VNIs.

Queue group redirection provides a more advanced mechanism where individual VNIs can use their own FP access ingress and port egress queue group instances for QoS control. Packets with a specific VNI can be redirected to an access ingress and egress queue group instance, with the policer or queue within the queue group instance being selected using the classification capabilities in the associated SAP ingress or egress QoS policy. Queue group redirection is enabled by applying a queue group redirect list, which contains VNI match statements, under an ingress and egress SAP that is configured with a default queue group instance.

VXLAN VNI queue group redirection is not supported at ingress on the 7750 SR-a4/a8 (which do not support FP ingress queue groups). It is supported under IES and VPRN interface SAPs, but not on an HSMDA. Redirecting fragmented packets based on the VNI to a queue group instance is not supported; the VNI matching is ignored and the default queue group instance is used.

7.2.5.1 Queue Group Redirect List

A queue group redirection list allows packets that match a configured type value to be redirected to a specific queue group instance at SAP ingress or egress.

The queue group instance must be an instance of the default queue group, which is the queue group configured under the SAP with the QoS policy.

The only type matching supported is **vxlان-vni**, which enables an ingress and egress match based on the VNI in IPv4 and IPv6 VXLAN and VXLAN GPE packets. Each entry in the redirect list matches a specific VNI and maps it to a specific queue group instance. A maximum of 16 match statements can be configured in a queue group redirect list.

A redirect list is configured as follows:

```
configure
  qos
    queue-group-redirect-list redirect-list-name
    [create]
      type redirect-list-type
      match field-value instance instance-id
    exit
```

where:

- *redirect-list-name* specifies the name of the queue group redirect list
- *redirect-list-type* specifies the type of queue group redirect list; the default and only possible value is **vxlان-vni**

- *field-value* specifies the value of the field in the ingress or egress packet which, when matched, will cause the packet to be redirected to the specified queue group instance. The *field-value* is dependent on the setting of the type and must be a valid VXLAN VNI. The VNI can be specified in any of the available formats but is always shown in decimal.
- *instance-id* specifies the instance of the queue group template to which the VXLAN and VXLAN GPE packets are redirected. The packets can be redirected to the default instance, that is, the instance specified with the QoS policy.

A queue group redirection list is applied under the SAP ingress or egress as shown in the following example. The prior configuration of a default queue group instance (the queue group instance specified with the QoS policy under the SAP ingress or egress) is mandatory. Only a single list can be applied per-SAP ingress or egress.

```
configure
  service
    {ies | vprn} service-id
      interface ip-int-name
        sap sap-id
          ingress
            qos policy-id fp-redirect-group
            queue-group-name instance instance-id
            queue-group-redirect-list
            redirect-list-name
          egress
            qos policy-id port-redirect-
            group queue-group-name instance instance-id
            queue-group-redirect-list
            redirect-list-name
```

In order for packets to be redirected to the policer or queue in the queue group instance specified in a queue group redirect list, the following criteria must be met:

- The FC in the SAP ingress or egress QoS policy must be redirected to an FP or port redirect queue group.
- The VNI to be matched must be configured in a match statement in the queue group redirect list.
- The instance corresponding to the related match statement must exist on the FP ingress or port egress through which the traffic ingresses or egresses.

If any of these conditions are not met, the traffic will not be redirected to the queue group instance specified in the redirect list and will fall back to either using the SAP policers and queues or the default instance (the one configured with the QoS policy under the SAP), depending on which condition is not met; this applies to VXLAN and VXLAN GPE packets with VNIs not listed in the redirect list, and non-VXLAN and non-VXLAN GPE packets.

Modifying a queue group redirect list will cause the VNI-to-instance matching to change on all SAPs (ingress and egress) where the queue group redirect list is applied.

There is a limit on the number of SAP/VNI combinations at ingress and the number of SAP-instance/VNI combinations at egress per FP, which can be seen using the **tools dump resource-usage card slot-num fp fp-number** command. The egress limit is a subset of the Dynamic Service Entries and instance mismatches are considered to be allocated on the related ingress and egress FP. The output shows the total, allocated, and free resources, as shown in the following example:

```
*A:PE# tools dump resource-usage card 1 fp 1
=====
Resource Usage Information for Card Slot #1 FP #1
=====
Total    Allocated    Free
-----
...
      Sap IngQGrp RedirLst Entries |      31999          2      31997
      Dynamic Service Entries +    131071          2    131069
      SapInst EgrQGrp RedirLst Entries -    31999          2      31997
...
=====
*A:PE#
```

If a queue group instance specified in an applied queue group redirect list does not exist on the SAP's FP ingress or port egress, a flag is displayed in the show output for the SAP, for example:

```
*A:PE# show service id 1 sap 1/1/1
=====
Service Access Points(SAP)
=====
Service Id      : 1
SAP             : 1/1/1                      Encap           : null
Description     : (Not Specified)
Admin State     : Up                        Oper State        : Up
Flags           : SapIngQGrpRedirMismatch
                  SapEgrQGrpRedirMismatch
Multi Svc Site  : None
Last Status Change : 01/31/2017 17:50:56
Last Mgmt Change  : 02/02/2017 12:02:56
=====
*A:PE#
```

The following show command can be used to identify which instances from the redirection list exist on the SAP's FP ingress and port egress and which do not exist and are causing the mismatch:

```
*A:PE# show service id 1 sap 1/1/1 queue-group-redirection
=====
Queue Group Redirect List Mismatch Information (Ingress SAP)
=====
```

```

-----
Queue Group      Instance      FP
-----
qg2              1          1/1
                2          1/1 : mismatch
-----
=====

Queue Group Redirect List Mismatch Information (Egress SAP)
=====
-----
Queue Group      Instance      Port
-----
qg2              1          1/1/1
                2          1/1/1 : mismatch
-----
=====
*A:PE#

```

7.2.5.2 Queue Group Redirect List Example

This section provides a configuration example for the use of queue group redirection lists.

In this example, each ingress queue group instance contains three policers and each egress queue group instance contains three queues (the details of the queue group templates and instance creation are not shown).

The VXLAN VNI mapping to ingress and egress queue group instances is as follows:

- the ingress and egress default mapping is to queue group qg1 instance 1
- VXLAN VNI 1 maps to ingress and egress instance 1 (the default instance)
- VXLAN VNI 2 maps to ingress and egress instance 2
- VXLAN VNI 3 maps to ingress and egress instance 3

Non-VXLAN traffic and VXLAN traffic with VNIs other than 2 and 3 use the default instance (VNI 1 is explicitly using the default instance).

An IP interface within an IES service is used.

QoS policies:

```

sap-ingress 10 create
    queue 1 create
    exit
    queue 11 multipoint create
    exit

```



```
        fc "af" create
            policer 2 fp-redirect-group
        exit
        fc "be" create
            policer 1 fp-redirect-group
        exit
        fc "ef" create
            policer 3 fp-redirect-group
        exit
        dscp af11 fc "af"
        dscp ef fc "ef"
    exit
sap-egress 10 create
    queue 1 create
    exit
    fc af create
        queue 2 port-redirect-group-queue
    exit
    fc be create
        queue 1 port-redirect-group-queue
    exit
    fc ef create
        queue 3 port-redirect-group-queue
    exit
    dscp af11 fc "af"
    dscp ef fc "ef"
exit
```

Queue group redirect list:

```
queue-group-redirect-list list1 create
    type vxlan-vni
    match 1 instance 1
    match 2 instance 2
    match 3 instance 3
exit
```

Service:

```
ies 1 customer 1 create
    interface "int" create
        address 10.0.0.2/30
        sap 1/1/1 create
            ingress
                qos 10 fp-redirect-group "qg1" instance 1
                queue-group-redirect-list "list1"
            exit
            egress
                qos 10 port-redirect-group "qg1" instance 1
```

```
                queue-group-redirect-list "list1"
            exit
        exit
    exit
exit
```

Show output:

```
A:PE# show qos queue-group-redirect-list detail
=====
Queue Group Redirect List Information
=====
Policy Name      : list1
Match Type      : vxlan-vni
-----
Match            Instance
-----
1                1
2                2
3                3
-----
Service Associations
-----
Service ID      Service Type      SAP
-----
1               IES               1/1/1      (Ingress/Egress)
-----
=====
A:PE#
```

7.3 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. The HSMDA also uses the term *queue group* to describe a group of eight preconfigured hardware queues on its egress port. When queue sharing and redirection is configured on egress, a set of eight HSMDA queues could be configured as a part of the queue group template. These correspond to eight hardware queues on the HSMDA. When all eight egress fcs are mapped to the queue-group instantiated in the egress port, the per-sap hsmda queue-group resource is freed.

7.4 Port Queue Groups

When 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.

Additional parameters can be configured under port queue groups, for example, an accounting policy, queue overrides, a scheduler policy, and scheduler overrides.

7.4.1 Percent-Rate Support

The **percent-rate** command is supported in an egress queue group template 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.

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]
```

7.5 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, or 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

When 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, or 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 IOMs, IMM, or XMA with ports in a LAG.

7.6 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.

7.7 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, the queues in a template can be created 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. (This step and the following step can be done in the opposite order.) 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. 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, the queues in a template can be created 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. (This step and the following step can be done in the opposite order.) 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*

Redirection to a queue group on the HSMDA supports the SAP-based provisioning model, only.

7.7.1 Ingress and Egress SAP Forwarding Class Redirection Association Rules

7.7.1.1 Policy Based Provisioning Model

The association rules between SAP ingress and egress QoS policies and queue group templates are as follows: 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 preceding operation is successful:

- 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 preceding operation is successful:

- 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 preceding operation is successful:

- 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.

7.7.1.2 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 preceding operation is successful:

- 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 preceding operation is successful:

- 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, 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.
- 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.

7.7.2 Access Queue Group Statistics

7.7.2.1 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 into a billing file based on the accounting policy applied to the queue group.

7.7.2.2 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.

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 into a billing file based on the accounting policy applied to the queue group.

7.8 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.

7.8.1 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 preceding operation is successful:

- 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 preceding operation is successful:

- The system finds all IP interfaces where the policy is applied.
- The system 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 preceding operation is successful:

- 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

7.8.2 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.

7.9 PW Shaping

7.9.1 Ingress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection

7.9.1.1 Feature Configuration

The user applies a network QoS policy to the ingress context of a spoke-SDP to redirect the mapping of a Forwarding Class (FC) to a policer defined in a queue-group template that is instantiated on the ingress Forwarding Plane (FP) where the PW packets are received. (This feature applies to both spoke-SDP and mesh-SDP. Spoke-SDP is used throughout for ease of reading.)

```
config>service>vprn>if>spoke-sdp>ingress>qos network-policy-id fp-redirect-group queue-group-name instance instance-id
```

Here, 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 that 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 an 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 broadcast-policer policer-id  
config>qos>network>ingress>fc>fp-redirect-group unknown-policer policer-id  
config>qos>network>ingress>fc>fp-redirect-group mcast-policer policer-id
```

The user can redirect the unicast, broadcast, unknown, and multicast packets of an FC to different policers to allow for different policing rates for these packet types (broadcast and unknown are only applicable to VPLS services). However, the queue-group is explicitly named only at the time the network QoS policy is applied to the spoke-SDP, as shown in the preceding 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:

config>qos>shared-queue policer-output-queues

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 including the PW template.

7.9.1.2 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 that needs to be redirected and optionally, for each traffic type (unicast, broadcast, unknown, or multicast).
2. Apply the queue-group template to the network ingress context of all IOM/IMM FPs 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.
 - 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 IOM/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 fails.
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 IOM/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 IOM/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 fails at the time the user associates the ingress context of a spoke-SDP with the named queue-group. In such a case, the PW packet feeds directly into 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 fails at the time the user associates the ingress context of a spoke-SDP with the named queue-group. In such a case, the PW packet feeds directly into 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 that 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 fed directly into 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, that is not explicitly redirected in the network QoS policy has the corresponding packets feed directly into 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 of whether 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 an 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.

7.9.1.3 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 the 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 of whether an instance of the named policer queue-group exists on the ingress FP that the PW packet is received on. The user can apply a QoS filter matching the dot1p 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 to which PW is redirected does not exist, the redirection command fails. In this case, the packet classification is performed according to the default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface that the PW packet is received on.

7.9.2 Egress PW Shaping using Spoke-SDP Forwarding Class-Based Redirection

7.9.2.1 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>if>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 an 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 in the preceding 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 an FC to use a queue in a queue-group, in which case there are no policers used.

```
config>qos>network>egress>fc>port-redirect-group queue queue-id
```
- The user can redirect an 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 an 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 including the PW template.

7.9.2.2 Provisioning Model

This provisioning model allows the mapping of one or more PWs to the same instance of queues, or policers and queue, 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 that needs to be redirected.
2. Apply the queue-group template to the network egress context of all IOM/IMM ports where there exists a network IP interface that 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-sdps 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 IOM/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 fails at the time the user associates the egress context of a spoke-SDP with 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 that 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 fails at the time the user associates the egress context of a spoke-SDP with 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 that 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 that 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 that is not explicitly redirected in the network QoS policy has the corresponding packets feed directly into the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

7.9.2.3 Egress Marking of PW Packet Header

When the queue-group name that the PW is redirected to exists and the redirection succeeds, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/dot1p/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 that 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/dot1p and the tunnel's DEI/dot1p/EXP and the DSCP or IP precedence will be remarked if **enable-dscp-prec-remarking** is enabled under the policer.

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 a non-existent queue-group name, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/dot1p/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 that the PW packet is forwarded to.

7.9.2.4 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.

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 do, however remain separate as one or more PWs can have FCs redirected to the same policer ID in the queue-group instance.

7.10 Queue Group Behavior on LAG

7.10.1 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.

7.10.2 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 100 Mb/s is defined on a queue group queue, each instance of the queue group (on each LAG port) will have a rate of 100 Mb/s.

7.10.3 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 defined 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.

7.10.4 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.

7.10.5 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.

7.11 Basic Configurations

7.11.1 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, explicitly map all forwarding classes to queue group queues. This rule is applicable to SAP ingress, SAP egress, and network QoS policies.

7.11.2 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#

```

7.11.3 Applying Ingress Queue Group to SAP Ingress Policy

The following displays 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#

```

7.11.4 Applying Egress Queue Group to SAP Egress Policy

The following displays 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#
```

7.11.5 Configuring SAP-based Egress Queue Redirection

The following displays a SAP egress policy configuration with port-redirect-group-queue construct (shown for both regular and HSMDA 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

```


7.11.6 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#

```

7.11.7 Configuring Overrides

The following output displays 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
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#

```

7.11.8 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 then egress port (or ports) on then 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 redirection: policy-based redirection and SAP-based redirection. 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 redirection 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 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 that 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.

7.11.9 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.

When a queue within a template is mapped by a forwarding class on any object, the queue may be edited, but not deleted.

7.11.10 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.

7.11.11 Using Queue Groups to Police Ingress/Egress Traffic on Network Interface

An example of the provisioning steps involved in using a queue-group to police ingress and egress traffic on a network interface is as follows:

```

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

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
    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

```

7.11.12 Configuring Ingress/Egress PW Shaping Using Spoke-SDP Forwarding Class-Based Redirection

An example of the provisioning steps involved in configuring PW shaping using spoke-SDP forwarding class-based redirection is as follows:

```

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
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
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
    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 198.51.100.0
        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-192.168.0.0"
        address 192.168.0.0/24
        port 3/2/1:1
    exit
    interface "system"
        address 192.168.0.1/32
    exit
#-----
```

7.11.13 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#
```

7.12 QoS Queue Group Template Command Reference

7.12.1 Command Hierarchies

- [Configuring Egress Queue Group Templates](#)
- [Configuring Ingress Queue Group Templates](#)
- [Configuring VXLAN VNI Queue Group Redirection](#)
- [Show Commands](#)
- [Monitor Commands](#)

7.12.1.1 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
          — fc fc-name [create]
          — no fc fc-name
            — queue queue-id
            — no queue
          — hsmdda-queues
            — low-burst-max-class class
            — no low-burst-max-class
            — packet-byte-offset {add add-bytes | subtract sub-bytes}
            — no packet-byte-offset
            — [no] queue queue-id
              — adaptation-rule [pir {max | min | closest}]
              — no adaptation-rule
              — burst-limit size [bytes | kilobytes]
              — no burst-limit
              — mbs {size [kilobytes | bytes] | default}
              — no mbs
              — rate pir-rate
              — no rate
              — slope-policy hsmdda-slope-policy-name
              — no slope-policy
              — wrr-weight weight
              — no wrr-weight
```

- **wrr-policy** *wrr-policy-name*
- **no wrr-policy**
- **policer** *policer-id* [create]
- **no policer** *policer-id*
 - **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**
 - [no] **enable-exceed-pir**
 - **high-prio-only** *percent-of-mbs*
 - **no high-prio-only**
 - **mbs** {size [bytes | kilobytes] | default}
 - **no mbs**
 - **packet-byte-offset** {add *add-bytes* | subtract *sub-bytes*}
 - **no packet-byte-offset**
 - **parent** {root | *arbiter-name*} [level *level*] [weight *weight-within-level*]
 - **no parent**
 - [no] **profile-capped**
 - **rate** {max | *pir-rate*} [cir {max | *cir-rate*}]
 - **no rate**
 - **stat-mode** {no-stats | minimal | offered-profile-no-cir | offered-profile-cir | offered-total-cir | offered-limited-capped-cir | offered-profile-capped-cir | offered-total-cir-exceed | offered-four-profile-no-cir | offered-total-cir-four-profile}
 - **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** {size [bytes | kilobytes] | default}
 - **no burst-limit**
 - **cbs** {size-in-kbytes | default}
 - **no cbs**
 - **drop-tail**
 - **exceed**
 - **percent-reduction-from-mbs** *percent*
 - **no percent-reduction-from-mbs**
 - **high**
 - **percent-reduction-from-mbs** *percent*
 - **no percent-reduction-from-mbs**
 - **highplus**
 - **percent-reduction-from-mbs** *percent*
 - **no percent-reduction-from-mbs**
 - **low**
 - **percent-reduction-from-mbs** *percent*

- **no percent-reduction-from-mbs**
- **[no] dynamic-mbs**
- **mbs** {size [bytes | kilobytes] | default}
- **no mbs**
- **packet-byte-offset** {add add-bytes | subtract sub-bytes}
- **no packet-byte-offset**
- **parent** scheduler-name [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]
- **no parent**
- **percent-rate** pir-percent [cir cir-percent]
- **no percent-rate**
- **pool** pool-name
- **no pool**
- **port-parent** [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]
- **no port-parent**
- **queue-delay** delay
- **no queue-delay**
- **rate** pir-rate [cir cir-rate]
- **no rate**
- **wred-queue** [policy slope-policy-name] [mode mode] [slope-usage slope-usage]
- **no wred-queue**
- **[no] queues-hqos-manageable**

7.12.1.2 Configuring Ingress Queue Group Templates

- ```

config
 — qos
 — queue-group-templates
 — ingress
 — 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 adaptation-rule] [cir adaptation-rule]
 — 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 add-bytes | subtract sub-bytes}
 — no packet-byte-offset

```

- **parent** {**root** | *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]
- **no parent**
- [no] **profile-capped**
- **rate** {**max** | *pir-rate*} [**cir** {**max** | *cir-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*] [**create**]
- **no queue** *queue-id*
  - **adaptation-rule** [**pir** *adaptation-rule*] [**cir** {**max** | **min** | **closest**}] [**fir** {**max** | **min** | **closest**}]
  - **no adaptation-rule**
  - **adv-config-policy** *adv-config-policy-name*
  - **no adv-config-policy**
  - **burst-limit** {*size* [**bytes** | **kilobytes**] | **default**}
  - **no burst-limit**
  - **cbs** {*size-in-kbytes* | **default**}
  - **no cbs**
  - [no] **cir-non-profiling**
  - **drop-tail**
    - **low**
      - **percent-reduction-from-mbs** *percent*
      - **no percent-reduction-from-mbs**
  - **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
  - **no mbs**
  - **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
  - **no packet-byte-offset**
  - **parent** *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]
  - **no parent**
  - **pool** *pool-name*
  - **no pool**
  - **rate** *pir-rate* [**cir** *cir-rate*] [**fir** *fir-rate*]
  - **rate** *pir-rate* **police**
  - **no rate**

### 7.12.1.3 Configuring VXLAN VNI Queue Group Redirection

- ```

config
  — qos
    — queue-group-redirect-list redirect-list-name [create]
    — no queue-group-redirect-list redirect-list-name
      — match field-value instance instance-id
      — no match field-value
      — type redirect-list-type
      — no type

```

7.12.1.4 Show Commands

show

- qos
 - **queue-group** [queue-group-name] {ingress | egress} [association | detail]
 - **queue-group** summary
 - **queue-group-redirect-list** [policy-name] [association | detail]

show

- **pools** mda-id[/port] [access-app [pool-name | service service-id | queue-group queue-group-name [instance instance-id]]]
- **pools** mda-id[/port] [network-app [pool-name | queue-group queue-group-name [instance instance-id]]]
- **pools** mda-id[/port] [direction [pool-name | service service-id | queue-group queue-group-name [instance id]]]
- **port** queue-group [{ingress | egress}] [queue-group-name] [{access | network}] [instance instance-id] [{statistics | associations | queue-depth [queue queue-id]]]
- **port** queue-group summary

7.12.1.5 Monitor Commands

For more information about monitor commands, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide* for command usage and CLI syntax.

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]]
- qos
 - arbiter-stats
 - **card** slot-number fp fp-number queue-group queue-group-name instance instance-id [ingress] [access | network] [interval seconds] [repeat repeat] [absolute | percent-rate [reference-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]
 - **port** port-id vport name [interval seconds] [repeat repeat] [absolute | rate]

7.12.2 Command Descriptions

7.12.2.1 Configuration Commands

7.12.2.1.1 Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>qgrps>egr>queue-group config>qos>qgrps>ing>>qgrps>ing>queue-group config>qos>qgrps>egr>qgrp>policer config>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.
Parameters	<i>description-string</i> — The description character string. Allowed values are any string up to 80 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

7.12.2.1.2 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.

ingress

Syntax	ingress
Context	config>qos>queue-group-templates
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.

queue-group

Syntax	queue-group <i>queue-group-name</i> [create] no queue-group <i>queue-group-name</i>
Context	config>qos>qgrps>egress config>qos>qgrps>ingress
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 no form of this 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 <i>queue-group-name</i> does not exist, the command has no effect and does not return an error.
Parameters	<i>queue-group-name</i> — Specifies the name of the queue group template up to 32 characters. 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. 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 <i>policer-id</i> [create] no policer <i>policer-id</i>
Context	config>qos>qgrps>ing>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 eight (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.

When 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 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.

When 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 enters that new policer's context for possible parameter modification.

Values 1 to 32

adaptation-rule

Syntax **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
no adaptation-rule

Context config>qos>queue-group-templates>egress>queue-group>policer

Description This command defines the method used by the system to derive the operational CIR and PIR settings when the policer 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.

When configured on an egress HSQ queue group queue, the **cir** keywords are ignored. This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the adaptation rule is performed under the **hs-wrr-group** within the egress queue group template.

When a specific **adaptation-rule** is removed, the default constraints for **pir** and **cir** apply.

The **no** form of this command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy.

Default	adaptation-rule pir closest cir closest
Parameters	pir — Defines the constraints enforced when adapting the policer's PIR. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the policer. When the pir parameter is not specified, the default constraint applies. cir — Defines the constraints enforced when adapting the policer's CIR. The cir parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the policer. When the cir parameter is not specified, the default constraint applies. max — Specifies that the operational rate for the policer will be equal to or less than the requested rate. min — Specifies that the operational rate for the policer will be equal to or greater than the requested rate. closest — Specifies that the operational rate for the policer will be the rate closest to the requested rate.

adv-config-policy

Syntax	adv-config-policy <i>adv-config-policy-name</i> no adv-config-policy
Context	config>qos>qgrps>egr>qgrp>policer config>qos>qgrps>egr>qgrp>queue config>qos>qgrps>ing>qgrp>policer config>qos>qgrps>ing>qgrp>queue
Description	This command specifies the name of the advanced configuration policy to be applied with this policer.
Parameters	<i>adv-config-policy-name</i> — Specifies an existing advanced configuration policy up to 32 characters.

cbs

Syntax	cbs { <i>size-in-kbytes</i> default } no cbs
Context	config>qos>qgrps>egr>qgrp>policer config>qos>qgrps>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	<i>size-in-kbytes</i> — For the queues, the size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes 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 0 to 16777216 or default Minimum default value: 16 Mbytes when CIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured CBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max CIR capped to 3968 kbytes, with a minimum of 256 bytes

high-prio-only

Syntax	high-prio-only <i>percent-of-mbs</i> no high-prio-only
Context	config>qos>qgrps>egr>qgrp>policer config>qos>qgrps>ing>qgrp>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.
Parameters	<i>percent-of-mbs</i> — The <i>percent-of-mbs</i> parameter is required when specifying high-prio-only and is expressed as a percentage. Values 0 to 100 Default 10

mbs

Syntax	mbs { <i>size</i> [bytes kilobytes] default } no mbs

Context	<pre>config>qos>qgrps>egr>qgrp>policer config>qos>qgrps>ing>qgrp>policer</pre>
Description	This command specifies the default maximum buffer size for the template queue in bytes or kilobytes. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. When 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>queue-group and port>ethernet>access>egress>queue-group contexts 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 that 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 packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard against 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. When configured on an egress queue group queue, this command and the queue-delay command are mutually exclusive. In order to change between the mbs and queue-delay parameters, the current parameter must be removed before adding the new parameter; that is, changing from mbs to queue-delay requires a no mbs before the queue-delay is configured and changing from queue-delay to mbs requires a no queue-delay before the mbs is configured. If queue-delay is configured for an egress queue group queue, it is not possible to override the MBS for that queue. 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. At ingress, trusted in-profile packets and untrusted high-priority packets use the policer's high-priority violate threshold while trusted out-of-profile and untrusted low-priority packets use the policer's low-priority violate threshold. At egress, in-plus-profile and in-profile packets use the policer's high-priority violate threshold and out-of-profile packets use the policer's low-priority violate threshold. Exceed-profile packets are discarded unless enable-exceed-pir is configured, in which case they are forwarded.

The PIR bucket's violate threshold represents 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	default
Parameters	size — 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 kbytes, 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. The queue MBS maximum value used is constrained by the pool size in which the queue exists and by the shared pool space in the corresponding megapool. Values 0 to 16777216 Default value: 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicitly configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes. [bytes kilobytes] — Select bytes or kilobytes. Default kilobytes default — Sets the MBS to its default value.

packet-byte-offset

Syntax	packet-byte-offset { add <i>add-bytes</i> subtract <i>sub-bytes</i> }
	no packet-byte-offset
Context	config>qos>qgrps>ing>qgrp>policer
Description	This command configures a packet byte offset for the QoS ingress queue-group policer.
Default	no packet-byte-offset
Parameters	add-bytes — Specifies the number of bytes to add as the offset amount. Values 0 to 31

sub-bytes — Specifies the number of bytes to add as the offset amount.

Values 1 to 32

parent

Syntax	parent { root <i>arbiter-name</i> } [level <i>level</i>] [weight <i>weight-within-level</i>] no parent
Context	config>qos>qgrps>egr>qgrp>policer config>qos>qgrps>ing>qgrp>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 or multiservice site context. To create a policer control hierarchy on an ingress or egress SAP context, a **policer-control-policy** must be applied to the SAP. When 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 is cleared and a resulting trap is 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** that 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. This functionality is available only for the 7450 ESS and 7750 SR.

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 or multiservice site 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

When 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.

Default **root**

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.

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

profile-capped

Syntax	[no] profile-capped
Context	config>qos>qgrps>egr>qgrp>policer config>qos>qgrps>ing>qgrp>policer
Description	This command enables a limit on the profile.
Default	no profile-capped

rate

Syntax	rate { max <i>pir-rate</i> } [cir { max <i>cir-rate</i> }] no rate
Context	config>qos>qgrps>egr>qgrp>policer config>qos>qgrps>ing>qgrp>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 each packet's size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches its 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 nor 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 kb/s (all packets out-of-profile). The no form of this command is used to restore the default metering and profiling rate to a policer.
Parameters	{max pir-rate} — Specifying the keyword max or an explicit <i>pir-rate</i> 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 <i>pir-rate</i> 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, 1 to 2000000000

cir {max | *cir-rate*} — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword **max** or an explicit *cir-rate* 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 kb/s. The *cir-rate* 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 CIR used is equivalent to **max**.

Values **max**, 0 to 2000000000

stat-mode

Syntax	stat-mode {no-stats minimal offered-profile-no-cir offered-priority-no-cir offered-profile-cir offered-priority-cir offered-total-cir offered-limited-profile-cir offered-profile-capped-cir offered-limited-capped-cir} no stat mode
Context	config>qos>qgrps>ing>qgrp>policer
Description	This command is used to configure the forwarding plane counters that allow offered, forwarded, and dropped accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, uncolored, 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, for example, 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 that prevents any packet accounting, the use of the policer's **parent** command requires that the policer's **stat-mode** 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. When 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. The total/allocated/free stats can be viewed by using the **tools dump resource-usage card fp** 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 stat-modes are described in [Table 60](#).

Table 60 Stat Mode Descriptions

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
no-stats	0	None	None	
Minimal	1	Single counter entering policer	Single counter for dropped/forwarded exiting policer	
offered-profile-no-cir	2	In/out entering policer	In/out entering policer	Intended for when the policer does not change the profile of packets. Includes only in- and out-of-profile.
offered-priority-no-cir	2	High/low entering policer	High/low entering policer	Intended for when only packet priority stats are required.
offered-profile-cir	4	In/out/uncolored entering policer	In/out exiting policer	Intended for when the policer can change the profile of packets to in- and out-of-profile.
offered-priority-cir	4	High/low entering policer	In/out exiting policer	Intended for when packet priority entering the policer and profile exiting the policer is required.
offered-total-cir	2	Single counter entering policer	In/out exiting policer	

Table 60 Stat Mode Descriptions (Continued)

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
offered-limited-profile-cir	3	Out/uncolored entering policer	In/out exiting policer	Intended for when the policer can change the profile of packet to in- and out-of-profile. The information is limited compared to offered-profile-capped-cir with the benefit of using one less stat resource.
offered-profile-capped-cir	5	In/out/uncolored entering policer	In/out exiting policer	Intended for when the policer has profile-capped configured.
offered-limited-capped-cir	4	In/uncolored entering policer	In/out exiting policer	Intended for when the policer has profile-capped configured. The information is limited compared to offered-profile-capped-cir with the benefit of using one less stat resource.

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, dropped and forwarded 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, no statistics are generated.

minimal — Counter resource allocation: 1

This **stat-mode** provides the minimal accounting resource usage and counter information, and includes the total offered, dropped and forwarded packet and octet counters for traffic entering (offered) and exiting (dropped/forwarded) the policer.

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 in-profile or out-of-profile 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 displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 61](#).

Table 61 Ingress Accounting Statistics Collected in minimal stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. All	apd	AllPacketsDropped
	aod	AllOctetsDropped
For. All	apf	AllPacketsForwarded
	aof	AllOctetsForwarded

offered-profile-no-cir — Counter resource allocation: 2

This **stat-mode** provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering the policer.

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 premarked (and trusted) packets. It is expected that, in this instance, a CIR rate will not be defined since all packets are already premarked. This mode does not prevent the policer from receiving untrusted (color undefined) nor does it prevent the policer from being configured with a CIR rate.

This mode is intended to be used without profile-capped configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 62](#).

Table 62 Ingress Accounting Statistics Collected in offered-profile-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-priority-no-cir — Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the packet priority of traffic entering the policer.

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 untrusted 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 premarked in-profile or out-of-profile nor does it prevent the policer from being configured with a CIR rate.

This mode is intended to be used without profile-capped configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 63](#).

Table 63 Ingress Accounting Statistics Collected in offered-priority-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. HiPrio	hpo	HighPriorityPacketsOffered
	hoo	HighPriorityOctetsOffered
Off. LowPrio	lpo	LowPriorityPacketsOffered
	loo	LowPriorityOctetsOffered
Dro. HiPrio	hpd	HighPriorityPacketsDropped
	hod	HighPriorityOctetsDropped
Dro. LowPrio	lpd	LowPriorityPacketsDropped
	lod	LowPriorityOctetsDropped
For. HiPrio	hpf	HighPriorityPacketsForwarded
	hof	HighPriorityOctetsForwarded
For. LowPrio	lpf	LowPriorityPacketsForwarded
	lof	LowPriorityOctetsForwarded

offered-profile-cir — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard in/out and uncolored. The offered counters cover traffic explicitly profiled to in-profile, traffic explicitly profiled to out-of-profile, and traffic that has not been explicitly profiled at ingress (uncolored).

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 untrusted 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 untrusted markings. If trusted in-profile packets are not being received, the **offered-limited-profile-cir** stat-mode could be used instead, which has the benefit of using a reduced number of stat resources.

This mode is intended to be used without **profile-capped** configured within the policer as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 64](#).

Table 64 Ingress Accounting Statistics Collected in offered-profile-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-priority-cir — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the priority of traffic entering the policer and the profile exiting the policer.

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 untrusted 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.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 65](#).

Table 65 Ingress Accounting Statistics Collected in offered-priority-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. HiPrio	hpo	HighPriorityPacketsOffered
	hoo	HighPriorityOctetsOffered
Off. LowPrio	lpo	LowPriorityPacketsOffered
	loo	LowPriorityOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-total-cir — Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter.

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 untrusted 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 untrusted packets.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 66](#).

Table 66 **Ingress Accounting Statistics collected in offered-total-cir stat-mode**

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-limited-profile-cir — Counter resource allocation: 3

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard out and uncolored. The offered counters cover traffic explicitly profiled to out-of-profile and traffic that has not been explicitly profiled at ingress (uncolored). The traffic explicitly profiled to in-profile is counted with the uncolored traffic.

The **offered-limited-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 untrusted packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile packets. If trusted in-profile packets are not being received, the **offered-limited-profile-cir** is preferred over **offered-profile-cir** because it uses a reduced number of stat resources.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 67](#).

Table 67 Ingress Accounting Statistics Collected in offered-limited-profile-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-profile-capped-cir — Counter resource allocation: 5

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard in/out and uncolored. The offered counters cover traffic explicitly profiled to in-profile, traffic explicitly profiled to out-of-profile, and traffic that has not been explicitly profiled at ingress (uncolored).

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 is that one of the counting resources in the forwarding plane and traffic manager will not be used.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 68](#).

Table 68 **Ingress Accounting Statistics Collected in offered-profile-capped-cir stat-mode**

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-limited-capped-cir — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed resulting in the traffic entering the policer comprising of hard in/out and uncolored. The offered counters cover in-profile traffic and traffic that has not been explicitly profiled at ingress (uncolored). The traffic explicitly profiled to out-of-profile is counted with the uncolored traffic.

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 **soft in-profile** with **profile in** (InProf) and **profile out** (OutProf) with **soft-out-of-profile** (Uncolor) and eliminates the 'offered undefined' statistic. If trusted out-of-profile packets are not being received, the **offered-limited-capped-cir** is preferred over **offered-profile-capped-cir** because it uses a reduced number of stat resources.

This mode is intended to be used with **profile-capped** configured within the policer.

The impact of using **offered-limited-capped-cir** stat-mode while **profile-capped** mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in [Table 69](#).

Table 69 Ingress Accounting Statistics Collected in offered-limited-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

queue

Syntax	queue <i>queue-id</i> [multipoint] [<i>queue-type</i>] [<i>queue-mode</i>] [create] no queue <i>queue-id</i>
Context	config>qos>qgrps>ing>queue-group
Description	This command creates a queue for use in a queue group template. When 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 that 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. When 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 this 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.
Parameters	<i>queue-id</i> — This required parameter identifies the queue that will either be created or edited within the queue group template. Values 1 to 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. Default non-multipoint (unicast queue) <i>queue-type</i> — The queue types are mutually exclusive. 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. Default best-effort

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 subclasses 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 defined for the queue. Low-priority packets are discarded at the low-priority MBS threshold that is derived from applying the low drop-tail percentage to the queue's MBS.

create — Keyword used to create the queue ID instance.

adaptation-rule

Syntax	adaptation-rule [pir { max min closest }] [cir { max min closest }] [fir { max min closest }] no adaptation-rule
Context	config>qos>queue-group-templates>ingress>queue-group>queue
Description	This command defines the method used by the system to derive the operational FIR, CIR and PIR settings when the queue is provisioned in hardware. For the FIR, CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint. When a specific adaptation-rule is removed, the default constraints for pir, cir and fir apply. The no form of this command removes any explicitly defined constraints used to derive the operational FIR, CIR and PIR created by the application of the policy.
Default	adaptation-rule pir closest cir closest fir closest
Parameters	pir — Defines the constraints enforced when adapting the queue's PIR. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the pir parameter is not specified, the default constraint applies.

- cir** — Defines the constraints enforced when adapting the queue's CIR. 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.
- fir** — Defines the constraints enforced when adapting the queue's FIR. The **fir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **fir** parameter is not specified, the default constraint applies. FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.
- max** — Specifies that the operational rate for the queue will be equal to or less than the requested rate.
- min** — Specifies that the operational rate for the queue will be equal to or greater than the requested rate.
- closest** — Specifies that the operational rate for the queue will be the rate closest to the requested rate.

burst-limit

Syntax	burst-limit { <i>size</i> [bytes kilobytes] default} no burst-limit
Context	config>qos>qgrps>egr>qgrp>queue config>qos>qgrps>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	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) bytes — Specifies that the value given for size must be interpreted as the burst limit in bytes.

kilobytes — Specifies that the value given for size must be interpreted as the burst limit in kilobytes.

cbs

Syntax	cbs { <i>size-in-kbytes</i> default } no cbs
Context	config>qos>qgrps>egr>qgrp>queue config>qos>qgrps>ing>qgrp>queue
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	<i>size-in-kbytes</i> — For the queues, the size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes 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 0 to 1048576 or default Minimum configurable non-zero value: 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4 Minimum non-zero default value: maximum of 10 ms of CIR or 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4

cir-non-profiling

Syntax	[no] cir-non-profiling
Context	config>qos>queue-group-templates>ingress>queue-group>queue
Description	This command prevents the modification of the profile of a packet-dependent queue rate compared to its configured CIR. The CIR continues to be used to affect the scheduling priority of a queue. The cir-non-profiling and the queue police commands are mutually exclusive. cir-non-profiling is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware. cir-non-profiling should not be configured under an ingress queue group template queue associated with a LAG which spans FP4-based and FP2/FP3-based hardware as the resulting operation could be different depending on which hardware type the traffic ingresses.

drop-tail

Syntax	drop-tail
Context	config>qos>qgrps>ing>qgrp>queue config>qos>qgrps>egr>qgrp>queue
Description	This command enters the context to configure queue drop-tail parameters.

low

Syntax	low
Context	config>qos>qgrps>ing>qgrp>queue>drop-tail config>qos>qgrps>egr>qgrp>queue>drop-tail
Description	This command enters the context to configure the queue low drop-tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets will not be accepted into the queue and will be discarded.

percent-reduction-from-mbs

Syntax	percent-reduction-from-mbs <i>percent</i> no percent-reduction-from-mbs
Context	config>qos>qgrps>ing>qgrp>queue>drop-tail>low
Description	This command configures the ingress queue group queue low drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, the low drop tail will be at 420 kbytes. Out-of-profile packets will not be accepted into the queue and will be discarded if the queue depth is greater than this value.
Default	10%
Parameters	<i>percent</i> — Specifies the percentage reduction from the MBS for a queue drop tail. Values 0 to 100, default

mbs

Syntax	mbs { <i>size</i> [bytes kilobytes] default } no mbs
Context	config>qos>qgrps>egr>qgrp>queue config>qos>qgrps>ing>qgrp>queue

Description This command specifies the default maximum buffer size for the template queue in bytes or kilobytes.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. When 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>queue-group** and **port>ethernet>access>egress>queue-group** contexts 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 that 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 packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard against 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.

When configured on an egress queue group queue, this command and the **queue-delay** command are mutually exclusive. In order to change between the **mbs** and **queue-delay** parameters, the current parameter must be removed before adding the new parameter; that is, changing from **mbs** to **queue-delay** requires a **no mbs** before the **queue-delay** is configured and changing from **queue-delay** to **mbs** requires a **no queue-delay** before the **mbs** is configured. If **queue-delay** is configured for an egress queue group queue, it is not possible to override the MBS for that queue.

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.

At ingress, trusted in-profile packets and untrusted high-priority packets use the policer's high-priority violate threshold while trusted out-of-profile and untrusted low-priority packets use the policer's low-priority violate threshold.

At egress, inplus-profile and in-profile packets use the policer's high-priority violate threshold and out-of-profile packets use the policer's low-priority violate threshold. Exceed-profile packets are discarded unless **enable-exceed-pir** is configured, in which case they are forwarded.

The PIR bucket's violate threshold represents 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	default
Parameters	size — 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 kbytes, 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. The queue MBS maximum value used is constrained by the pool size in which the queue exists and by the shared pool space in the corresponding megapool. Values 0 to 1048576 or default Minimum configurable non-zero value: 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4 Minimum non-zero default value: maximum of 10 ms of CIR, or 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4 [bytes kilobytes] — Select bytes or kilobytes. Default kilobytes default — Sets the MBS to its default value.

packet-byte-offset

Syntax	packet-byte-offset { add <i>add-bytes</i> subtract <i>sub-bytes</i> } no packet-byte-offset
Context	config>qos>qgrps>ing>qgrp>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 ingress scheduling and profiling 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. Both the scheduling and profiling throughput is affected by the offset as well as the stats (accounting) associated with the queue. The packet-byte-offset does not apply to drop statistics, received valid statistics, or the offered managed and unmanaged statistics used by Ingress Multicast Path Management. 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 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 queue for scheduling, profiling and accounting purposes. From the queue's perspective, the packet size is increased by the amount being added to the size of each packet.

Values 0 to 30, in steps of 2

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 queue for scheduling, profiling and accounting purposes. From the queue's perspective, the packet size is reduced by the amount being subtracted from the size of each packet. The minimum resulting packet size used by the system is 1 byte.

Values 0 to 64, in steps of 2

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>qgrps>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, policers (at egress only), 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 multiservice 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 SAP that 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 and automatically returns 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 this 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. When 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, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

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 policers, queues, and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

All **weight** values from all weighted active queues, policers, 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, policer, or scheduler. 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 policers, queues, and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all non-zero weighted queues, policers, and schedulers at that level are operating at the maximum bandwidth or are idle.

Values 0 to 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*. Queues or schedulers will not receive parental bandwidth until all 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 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 relative to their weights.

Values 1 to 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 to 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 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 to 8 (8 is the highest priority)

Default 0

pool

Syntax **pool** *pool-name*
no pool

Context config>qos>qgrps>egr>qgrp>queue
config>qos>qgrps>ing>qgrp>queue

Description This command 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. Named pools are not supported on the 7950 XSR, 7750 SRc4/12, 7750 SR-a4/a8, or 7750 SR-1e/2e/3e.

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. When the pool comes into existence on the MDA or port, the queue will be mapped to the new pool.

When 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 on the **queue** command output using the **pool** keyword if it is configured.

The **no** form of this command removes the pool name from the configuration.

Parameters *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. The specified name must be an ASCII name string up to 16 characters.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*] [**fir** *fir-rate*]
rate *pir-rate* **police**
no rate

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

Description This command defines the administrative Peak Information Rate (PIR), the administrative Committed Information Rate (CIR), and the administrative Fair Information Rate (FIR) 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, unless **cir-non-profiling** is configured. In-profile, then out-of-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 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 FIR defines an additional rate at which the system prioritizes the queue over other queues competing for the same bandwidth above that used by the CIR.

The **rate** command can be executed at any time, altering the PIR, CIR, and FIR for all queues created through the association of the ingress queue group template with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR, CIR, and FIR parameters (**max**, 0, 0).

Default rate max cir 0 fir 0

- Parameters** *pir-rate* — Defines the administrative PIR, 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 the value must be given as a positive integer. The actual PIR is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.
- Values** 1 to 2000000000 kb/s, **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 the value must be given as a positive integer. The actual CIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.
- Values** 1 to 2000000000 kb/s, **max**
- Default** 0
- fir-rate* — The **fir** parameter overrides the default administrative FIR used by the queue. When the rate command is executed, an FIR setting is optional. When the rate command has not been executed or the **fir** parameter is not explicitly specified, the default FIR (0) is assumed. Fractional values are not allowed and the value must be given as a positive integer. The actual FIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned. FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.
- Values** 1 to 2000000000 kb/s, **max**
- Default** 0
- police** — Specifies that traffic feeding into the queue instance above the specified rate is dropped.

egress

- Syntax** **egress**
- Context** config>qos>queue-group-templates
- Description** This command enters 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.

fc

Syntax	fc <i>fc-name</i> [create] no fc <i>fc-name</i>
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 <i>fc-name</i>. 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. On the 7450 ESS and 7750 SR, egress policed subscriber traffic works in a slightly different way. The subscriber and subscriber host support destination and organization strings 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	<i>fc-name</i> — A valid forwarding class must be specified as <i>fc-name</i> when the fc command is executed. When the fc <i>fc-name</i> command is successfully executed, the system will enter the specified forwarding class context where the queue <i>queue-id</i> command may be executed. Values be, l1, af, l2, h1, ef, h2, nc

queue

Syntax	queue <i>queue-id</i>

no queue

Context	config>qos>qgrps>egr>queue-group-template>fc
Description	This command is used to map the forwarding class to the specified <i>queue-id</i> . The specified <i>queue-id</i> must exist within the egress queue group template. When 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	<i>queue-id</i> — The specified <i>queue-id</i> must exist within the egress queue group template.
Values	1 to 8

hsmdda-queues

Syntax	hsmdda-queues
Context	config>qos>qgrps>egr>qgrp
Description	This command enables the context to configure HSMDDA queues.

low-burst-max-class

Syntax	low-burst-max-class <i>class</i> no low-burst-max-class
Context	config>qos>qgrps>egr>qgrp>hsmdda-queues
Description	This command assigns the low burst maximum class to associate with the HSMDDA queue. The no form of this command is used to remove per packet size modifications from the queue.
Parameters	<i>class</i> — Specifies the low burst max class
Values	1 to 8

packet-byte-offset

Syntax	packet-byte-offset { add <i>add-bytes</i> subtract <i>sub-bytes</i> } no packet-byte-offset
Context	config>qos>qgrps>egr>qgrp>hsmdda-queues config>qos>qgrps>egr>qgrp>policer config>qos>qgrps>egr>qgrp>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 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 or policer instance, it adjusts the immediate packet size. This means that the queue rates (i.e., operational PIR and CIR) and policer or queue bucket updates use the adjusted packet size. In addition, the 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 applies for the agg-rate-limit on a SAP, a subscriber, or a Multiservice 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 policer or queue rate will be capped to a user-configured on-the-wire rate and the packet-byte-offset is not included; however the offsets are applied to the statistics.

The **no** form of this command is used to remove per packet size modifications from the queue.

Parameters *add-bytes* — Specifies that 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 to 32

sub-bytes — Specifies that the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the queue for scheduling and accounting purposes. The minimum resulting packet size used by the system is 1 byte.

Values 0 to 64

queue

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

Context config>qos>qgrps>egr>qgrp>hsmdda-queues

Description This command is used to map the HSMDDA queue to the specified *queue-id*. The specified *queue-id* must exist within the egress queue group template.

Parameters *queue-id* — The specified *queue-id* must exist within the egress queue group template.

Values 1 to 8

adaptation-rule

Syntax	adaptation-rule [<i>pir adaptation-rule</i>] no adaptation-rule
Context	config>qos>qgrps>egr>qgrp>hsmda-queues
Description	This command specifies how the system should resolve differences between the specified scheduling limit derived from the WRR group's rate command and the actual operational rate obtainable in hardware. The min, max, and closest mutually exclusive keywords specify whether the next highest rate or next lowest rate should be selected by the system. The no form of the command reverts to the default value.
Default	adaptation-rule pir closest
Parameters	pir <i>adaptation-rule</i> — Defines the constraints enforced when adapting the policer's PIR. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the policer. When the pir parameter is not specified, the default constraint applies. Values max — The max option is mutually exclusive to the min and closest options. When max is specified, the operational PIR for the policer will be equal to or less than the requested rate. min — The min option is mutually exclusive to the max and closest options. When min is specified, the operational PIR for the policer will be equal to or greater than the requested rate. closest — The closest parameter is mutually exclusive to the min and max parameter. When closest is specified, the operational PIR for the policer will be the rate closest to the requested rate.

burst-limit

Syntax	burst-limit <i>size</i> [bytes kilobytes] no burst-limit
Context	config>qos>qgrps>egr>qgrp>hsmda-queues>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 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	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) bytes — Specifies that the value given for size must be interpreted as the burst limit in bytes. kilobytes — Specifies that the value given for size must be interpreted as the burst limit in kilobytes.

mbs

Syntax	mbs {size [bytes kilobytes] default} no mbs
Context	config>qos>qgrps>egr>qgrp>hsmda-queues>queue
Description	This command specifies the default maximum buffer size for the template queue in bytes or kilobytes.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. When 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>queue-group** and **port>ethernet>access>egress>queue-group** contexts 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 that 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 packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard against 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.

When configured on an egress queue group queue, this command and the **queue-delay** command are mutually exclusive. In order to change between the **mbs** and **queue-delay** parameters, the current parameter must be removed before adding the new parameter; that is, changing from **mbs** to **queue-delay** requires a **no mbs** before the **queue-delay** is configured and changing from **queue-delay** to **mbs** requires a **no queue-delay** before the **mbs** is configured. If **queue-delay** is configured for an egress queue group queue, it is not possible to override the MBS for that queue.

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.

At ingress, trusted in-profile packets and untrusted high-priority packets use the policer's high-priority violate threshold while trusted out-of-profile and untrusted low-priority packets use the policer's low-priority violate threshold.

At egress, inplus-profile and in-profile packets use the policer's high-priority violate threshold and out-of-profile packets use the policer's low-priority violate threshold. Exceed-profile packets are discarded unless **enable-exceed-pir** is configured, in which case they are forwarded.

The PIR bucket's violate threshold represents 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	default
Parameters	size — 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 kbytes, 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. The queue MBS maximum value used is constrained by the pool size in which the queue exists and by the shared pool space in the corresponding megapool. Values 0 to 2625 kilobytes, 0 to 2688000 bytes, or default [bytes kilobytes] — Select bytes or kilobytes. Default kilobytes default — Sets the MBS to its default value.

rate

Syntax	rate <i>pir-rate</i> no rate
Context	config>qos>qgrps>egr>qgrp>hsmda-group>queue

Description	This command specifies the PIR rates. The no form of the command reverts to the default value.
Default	rate max
Parameters	<i>pir-rate</i> — Specifies the PIR rates in Kb/s. Values 1 to 1000000000, max

slope-policy

Syntax	slope-policy <i>hsmda-slope-policy-name</i> no slope-policy
Context	config>qos>qgrps>egr>qgrp>hsmda-group>queue
Description	
Parameters	<i>hsmda-slope-policy-name</i> —

wrr-weight

Syntax	wrr-weight <i>weight</i> no wrr-weight
Context	config>qos>qgrps>egr>qgrp>hsmda-queues>queue
Description	This command specifies the WRR weight which this queue should parent into the scheduler. The weight of each queue determines how much bandwidth that queue gets out of the total rate for the scheduling class. The no form of the command reverts to the default.
Default	wrr-weight 1
Parameters	<i>weight</i> — Specifies WRR weight. Values 1 to 32

wrr-policy

Syntax	wrr-policy <i>wrr-policy-name</i> no wrr-policy
Context	config>qos>qgrps>egr>qgrp>hsmda-queues
Description	This command configures an HSM DA egress WRR policy.

The **no** form of the command removes the policy from the configuration.

Parameters *wrr-policy-name* — Specifies the WRR policy name, up to 32 characters.

policer

Syntax **policer** *policer-id* [**create**]
no policer *policer-id*

Context *cfg>qos>qgrps>egr>queue-group*

Description This command configures a QoS egress 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 to 8

create — This optional keyword creates an ingress queue-group policer.

enable-exceed-pir

Syntax [**no**] **enable-exceed-pir**

Context *cfg>qos>qgrps>egr>qgrp*

Description This command enables the forwarding of traffic exceeding the PIR for a SAP egress or a network egress queue group (configured in the egress queue group template) policer. This traffic is forwarded as exceed-profile instead of being dropped.

Default no enable-exceed-pir

adaptation-rule

Syntax **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
no adaptation-rule

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

When configured on an egress HSQ queue group queue, the **cir** keywords are ignored. This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the adaptation rule is performed under the **hs-wrr-group** within the egress queue group template.

When a specific **adaptation-rule** is removed, the default constraints for **pir** and **cir** apply.

The **no** form of this command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy.

Default	adaptation-rule pir closest cir closest
Parameters	pir — Defines the constraints enforced when adapting the queue's PIR. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the pir parameter is not specified, the default applies. cir — Defines the constraints enforced when adapting the queue's CIR defined. 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 — Specifies that the operational rate for the queue will be equal to or less than the requested rate. min — Specifies that the operational rate for the queue will be equal to or greater than the requested rate. closest — Specifies that the operational rate for the queue will be the rate closest to the requested rate.

exceed

Syntax	exceed
Context	cfg>qos>qgrps>egr>qgrp>queue>drop-tail
Description	This command enters the context to configure the queue exceed drop-tail parameters. The exceed drop tail defines the queue depth beyond which exceed-profile packets will not be accepted into the queue and will be discarded.

percent-reduction-from-mbs

Syntax	percent-reduction-from-mbs <i>percent</i> no percent-reduction-from-mbs
Context	config>config>qos>qgrps>egr>qgrp>queue>drop-tail>exceed config>qos>qgrps>egr>qgrp>queue>drop-tail>high

	<pre>config>qos>qgrps>egr>qgrp>queue>drop-tail>highplus config>qos>qgrps>>qgrps>egr>egr>qgrp>queue>drop-tail>low</pre>
Description	This command configures the egress queue group queue drop tails as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, the low drop tail will be at 420 kbytes. Out-of-profile packets will not be accepted into the queue and will be discarded if the queue depth is greater than this value. The drop tails apply to packets with the following profile states: • exceed drop tail: exceed-profile • high drop tail: in-profile • highplus drop tail: inplus-profile • low drop tail: out-of-profile
Default	exceed drop tail: 20% low drop tail: 10% high drop tail: 0% highplus drop tail: 0%
Parameters	<i>percent</i> — Specifies the percentage reduction from the MBS for a queue drop tail. Values 0 to 100, default

high

Syntax	high
Context	cfg>qos>qgrps>egr>qgrp>queue>drop-tail
Description	This command enters the context to configure the queue high drop-tail parameters. The high drop tail defines the queue depth beyond which in-profile packets will not be accepted into the queue and will be discarded.

highplus

Syntax	highplus
Context	cfg>qos>qgrps>egr>qgrp>queue>drop-tail
Description	This command enters the context to configure the queue highplus drop-tail parameters. The highplus drop tail defines the queue depth beyond which inplus-profile packets will not be accepted into the queue and will be discarded.

dynamic-mbs

Syntax	[no] dynamic-mbs
Context	config>qos>qgrps>egr>qgrp>queue
Description	This command enables support for dynamically modifying the MBS size of a queue using HQoS in order to maintain the maximum latency for traffic in the queue based on the queue's configured MBS and the ratio of its operational PIR to its administrative PIR. As the HQoS algorithm updates the operational PIR, by reducing or increasing it, the MBS of the queue is adjusted accordingly. The configuration of dynamic MBS and the configuration of queue depth monitoring (monitor-depth command) are mutually exclusive. Queue depth monitoring is an override on the queue where the queue group is applied. The no form of this command disables dynamic MBS resizing.
Default	no dynamic-mbs

percent-rate

Syntax	percent-rate <i>pir-percent</i> [<i>cir cir-percent</i>] no percent-rate
Context	config>qos>qgrps>egr>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 <i>queue-id</i> will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gb and a 10-Gb Ethernet port, the queue's rates will be 10 times greater on the 10 Gb 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. When configured on an egress HSQ queue group queue, the cir keyword is ignored. This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case the configuration of the rate is performed under the hs-wrr-group within the egress queue group template.

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. Similarly, 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 any time.

The **no** form of this command returns the queue to its default shaping rate and CIR rate.

Parameters *pir-percent* — Expresses 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 0.01 to 100.00 percent

Default 100.0

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 change dynamically due to configuration or auto-negotiation and the egress rate setting.

Values 0.01 to 100.00 percent

Default 100.0

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>queue-group-template>queue

Description This command defines the port scheduling parameters used to control the queue's 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 and the **parent** command are mutually exclusive.

The **no** form of this 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 to 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 to 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 is used.

Values 0 to 100

Default 0

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 is used.

Values 0 to 8 (8 is the highest priority)

Default 0

queue-delay

Syntax **queue-delay** *delay*
no queue-delay

Context config>qos>qgrps>egr>qgrp>queue

Description This command configures the target queue delay for packets forwarded through the queue. It is used to determine the related queue parameters based on the administrative PIR of the queue. This command and the **mbs** command are mutually exclusive.

In order to change between the **mbs** and **queue-delay** parameters, the current parameter must be removed before adding the new parameter; that is, changing from **mbs** to **queue-delay** requires a **no mbs** before the **queue-delay** is configured and changing from **queue-delay** to **mbs** requires a **no queue-delay** before the **mbs** is configured.

If **queue-delay** is configured for an egress queue group queue, it is not possible to override the MBS for that queue.

The **no** form of this command disables the determination of the queue parameters based on the queue delay.

Default no queue-delay

Parameters *delay* — Specifies the target queue delay in ms.

Values 0 to 5000 (decimal)

rate

Syntax **rate** *pir-rate* [*cir cir-rate*]
no rate

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

Description This command defines the administrative PIR and the administrative CIR parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress port. 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 any time, altering the PIR and CIR for all queues created through the association of the egress queue group template with the *queue-id*.

When configured on an egress HSQ queue group queue, the **cir** keyword is ignored.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the rate is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this 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

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 to 200000000 kb/s, **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 to 200000000 kb/s, **max**

Default 0

wred-queue

Syntax	wred-queue [<i>policy slope-policy-name</i>] [mode { native pool-per-queue }] [slope-usage { default exceed-low }] no wred-queue
Context	config>qos>qgrps>egr>qgrp>queue
Description	This command allows the configuration of WRED per queue with the following options: • Native hardware WRED This uses the hardware per queue WRED capabilities of FP3- and higher-based hardware and is configured with the native keyword. • Pool per queue WRED This implements each queue in its own pool and uses the WRED capabilities of the pool to provide WRED per queue. This is configured with the pool-per-queue keyword.

Native Hardware WRED

When the **wred-queue mode native** command is configured, the queue uses the WRED capabilities of FP3- and higher-based hardware. In this case, the out and exceed-profile traffic map to the low and exceed WRED slopes specified within the slope policy, and the inplus and in-profile traffic uses the MBS drop tail; this requires the **slope-usage** to be configured as **exceed-low**. The instantaneous queue depth is compared against the low and exceed slopes so the time average factor in the slope policy is ignored.

When a policy is not explicitly defined, the default slope policy is used.

When **native** mode is enabled for a queue, the drop-tail commands are ignored; traffic mapped to a slope that is shut down will use the MBS drop tail.

Native hardware WRED is supported on FP3- and higher-based hardware and is ignored on FP2 hardware.

The **no** form of this command restores the queue default congestion control behavior to the queue.

Pool-per-queue WRED

When the **wred-queue mode pool-per-queue** command is defined and the queue ID is created, 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 are not changed.

In the case where this is applied with WRED queue support shut down (**config>card>fp>egress>wred-queue-control>shutdown**), the queue will continue to map to either its default pool or the pool defined in the **pool** command. If the **no shutdown** command is executed in the **wred-queue-control** context, the queue will 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 forwarding plane. 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 highplus, high, low, and exceed 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 the highplus, high, low, or exceed WRED slope based on the packet's profile. If the packet is inplus-profile, the highplus slope is used. If the packet is in-profile, the high slope is used. If the packet is out-of-profile, the low slope is used, and if the packet is exceed-profile, the exceed slope is used. This mapping of packet profile to slope is enabled using the **slope-usage default** parameter. Each WRED slope performs a probability discard based on the current weighted average pool depth.

When **wred-queue** is enabled for an egress queue group queue, the queue pool and drop-tail commands are ignored; traffic mapped to a slope that is shut down will use the MBS drop tail.

The resource usage for the wred-queue pool-per-queue per forwarding plane can be seen in the **tools dump resource-usage card [slot-num] fp [fp-number]** output under *Dynamic Q2 Wred Pools*.

The **no** form of this 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. The queue then uses the default congestion control behavior.

Default	no wred-queue
Parameters	<i>slope-policy-name</i> — Specifies an existing slope policy that is used to override the default WRED slope policy. mode {native pool-per-queue} — Specifies whether the WRED per queue is using the native FP3- and higher-based hardware WRED capabilities or pool per queue. Values native - Each queue uses the hardware per queue WRED capabilities of the FP3- and higher-based hardware and requires slope-usage exceed-low. pool-per-queue - Each queue uses its own pool and the WRED capabilities of the pool to provide WRED per queue. This requires slope-usage default. Default native slope-usage {default exceed-low} — Specifies the type of congestion control to be used. Values default - Maps the inplus-profile, in-profile, out-of-profile, and exceed-profile traffic to the highplus, high, low, and exceed WRED slopes, respectively; this is only supported for pool-per-queue mode. exceed-low - Maps the out-of-profile and exceed-profile traffic to the low and exceed WRED slopes, with the inplus-profile and in-profile traffic using the MBS drop tail. This option is only supported for native mode. Default exceed-low

queues-hqos-manageable

Syntax	[no] queues-hqos-manageable
Context	config>qos>qgrps>egr>qgrp
Description	This command specifies that the queues within this egress queue group template are to be managed by the Hierarchical QoS (HQoS) process when the template is applied to an Ethernet access egress or network egress context. It is applicable to all egress queue group templates, including the default policer-output-queues template.

The **no queues-hqos-manageable** command must be configured for access egress queue groups that are used for post-policer traffic in order to prevent HQoS from measuring the traffic through a policer managed by HQoS then again through a post-policer access egress queue group queue.

Avoid scenarios that result in traffic being either not being measured or being measured twice by HQoS as they will cause the HQoS result to be inaccurate.

A template configured for **no queues-hqos-manageable** cannot be applied to an Ethernet network egress context. Any egress queue group templates applied to an Ethernet network egress context cannot be configured as **no queues-hqos-manageable**. The configuration of **no queues-hqos-manageable** and the configuration of policers and queue **packet-byte-offset** within the egress queue group template are mutually exclusive.

When a queue group template with **no queues-hqos-manageable** is configured under a port's Ethernet access egress context, the configuration of an aggregate rate or scheduler policy is not permitted under that context, nor are parent overrides for any of the queues in the queue group. If a port scheduler is configured on the port, the queue group queues are not parented to the port scheduler.

The **no** form of this command specifies that queues within this egress queue group are not managed by HQoS.

Default queues-hqos-manageable

queue-group-redirect-list

Syntax	queue-group-redirect-list <i>redirect-list-name</i> [create] no queue-group-redirect-list <i>redirect-list-name</i>
Context	config>qos
Description	This command configures a queue group redirect list that is used to redirect traffic to different instances of a queue group. The no form of this command deletes the queue group redirect list. A list can only be deleted when there no references to it.
Parameters	<i>redirect-list-name</i> — Specifies the name of the queue group redirect list, up to 32 characters.

match

Syntax	match <i>field-value</i> instance <i>instance-id</i> no match <i>field-value</i>
Context	config>qos>queue-group-redirect-list
Description	This command configures the value of the field in the ingress or egress packet which, when matched, will cause the packet to be redirected to the specified queue group instance. The <i>field-value</i> is dependent on the setting of the type and therefore must be a valid VXLAN VNI. A maximum of 16 match statements are supported in a queue group redirect list.

The **no** form of this command removes the match statement from the redirect list.

Parameters	<i>field-value</i> — Specifies the value of the field in the ingress or egress packet which, when matched, will cause the packet to be redirected to the specified queue group instance. Because the only permitted type is vlan-vni, the field must be a valid VXLAN VNI. The VNI can be specified in any of the available formats but is always shown in decimal. Values 1 to 16777215 (Decimal) 0x1 to 0xFFFFFFFF (Hexadecimal) 0b1 to 0b11111111111111111111111111111111 (Binary) <i>instance-id</i> — Specifies the instance of the queue group template to which the VXLAN traffic is redirected. The traffic can be redirected to the default instance, which is the instance specified with the QoS policy under the SAP ingress or egress. Values 1 to 65535

type

Syntax	type <i>redirect-list-type</i> no type
Context	config>qos>queue-group-redirect-list
Description	This command configures the type of a queue group redirect list. The default value is vlan-vni , which is the only possible value.
Parameters	<i>redirect-list-type</i> — Specifies the type of the queue group redirect list. The queue group redirect list is used to match VXLAN VNIs in IPv4 and IPv6 VXLAN or VXLAN GPE packets. Values vlan-vni

stat-mode

Syntax	stat-mode {no-stats minimal offered-profile-no-cir offered-profile-cir offered-total-cir offered-limited-capped-cir offered-profile-capped-cir offered-total-cir-exceed offered-four-profile-no-cir offered-total-cir-four-profile} no stat mode
Context	cfg>qos>qgrps>egr>qgrp>policer
Description	The sap-egress QoS policy's policer stat-mode command is used to configure the forwarding plane counters that allow offered, forwarded, and dropped 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, out-of-profile, and exceed-profile due to egress profile overrides) 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, for example, will not be configured with a CIR profiling rate and not all policers will receive explicitly reprofiled 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 that prevents any packet accounting, the use of the policer's **parent** command requires that the policer's **stat-mode** 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. When 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. The total/allocated/free stats can be viewed by using the **tools dump resource-usage card fp** 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 ingress policer stat-modes are described in [Table 70](#).

Table 70 Egress Policer Stat Mode Summary

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
no-stats	0	None	None	
Minimal	1	Single counter entering policer	Single counter for dropped/forwarded exiting policer	
offered-profile-no-cir	2	In/out entering policer	In/out entering policer	Intended for when the policer does not change the profile of packets. Includes only in- and out-of-profile.
offered-profile-cir	4	In/out/uncolored (that corresponds to in- or out-of-profile from the ingress processing) entering policer	In/out exiting policer	Intended for when the policer can change the profile of packets to in- and out-of-profile.

Table 70 Egress Policer Stat Mode Summary (Continued)

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
offered-total-cir	2	Single counter entering policer	In/out exiting policer	
offered-limited-capped-cir	4	In/out entering policer	In/out exiting policer	Intended for when the policer has profile-capped configured. The information is limited compared to offered-profile-capped-cir with the benefit of using one less stat resources.
offered-profile-capped-cir	5	In/out/uncolored (that corresponds to in- or out-of-profile from the ingress processing) entering policer	In/out exiting policer	Intended for when the policer has profile-capped configured.
offered-total-cir-exceed	3	Single counter entering policer	In/out/exceed exiting policer	Intended for when the policer is configured with enable-exceed-pir to forward packets that exceed its configured PIR or when traffic is egress reclassified to profile exceed.
offered-four-profile-no-cir	4	Inplus/in/out/exceed entering policer	Inplus/in/out/exceed entering policer	Intended to be used when the policer does not change the profile of the packets and traffic is egress reclassified to profile inplus and/or exceed.

Table 70 Egress Policer Stat Mode Summary (Continued)

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
offered-total-cir-four-profile	4	Single counter entering policer	Inplus/in/out/exceed exiting policer	Intended to be used when the policer can change the profile of the packet and traffic is egress reclassified to profile inplus.

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, no statistics are generated.

minimal — Counter resource allocation: 1

This stat-mode provides the minimal accounting resource usage and counter information, and includes only the total offered, dropped and forwarded packet and octet counters for traffic entering (offered) and exiting (dropped/forwarded) the policer.

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates one 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 and do not count different profile output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate or using exceed PIR.

The counters displayed in the **show** output and those collected when **collect-stats** is enabled are described in [Table 71](#) (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 71 Egress Accounting Statistics Collected in minimal stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. All	apd	AllPacketsDropped
	aod	AllOctetsDropped
For. All	apf	AllPacketsForwarded
	aof	AllOctetsForwarded

offered-profile-no-cir — Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering the policer. inplus-profile traffic is counted with the in-profile counters and exceed-profile traffic is counted with the out-of-profile counters.

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, dropped and forwarded statistics are required from the egress policer, but a CIR or **enable-exceed-pir** 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 or using **enable-exceed-pir**.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in [Table 72](#) (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 72 Egress Accounting Statistics Collected in offered-profile-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered

Table 72 Egress Accounting Statistics Collected in offered-profile-no-cir stat-mode (Continued)

Show Output	Accounting Stats Collected	
	Field	Field Description
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-profile-cir — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer comprises of hard inplus/in/out/exceed and soft in/out. The offered counters cover traffic reclassified to in-profile (that includes traffic reclassified to inplus-profile), traffic reclassified to out-of-profile (that includes traffic reclassified to exceed-profile), and traffic that has not been reclassified at egress (Uncolor). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

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 profile-based offered, dropped and forwarded 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.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in [Table 73](#) (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 73 Egress Accounting Statistics Collected in offered-profile-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-total-cir — Counter resource allocation: 2

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. In the dropped and forwarded counters, in-plus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

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 untrusted 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 untrusted packets.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in [Table 74](#) (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 74 Egress Accounting Statistics Collected in offered-total-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-limited-capped-cir — Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer comprises of hard inplus/in/out/exceed and soft in/out. The offered counters cover in-profile traffic (that includes traffic reclassified to inplus-profile) and out-of-profile traffic (that includes traffic reclassified to exceed-profile). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

When **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 **soft-in-profile** with **profile in** and **soft-out-of-profile** with **profile out** and eliminates the offered-undefined statistic.

The impact of using **offered-limited-capped-cir** stat-mode while profile-capped mode is disabled is 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.

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in [Table 75](#) (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 75 Egress Accounting Statistics Collected in offered-limited-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-profile-capped-cir — Counter resource allocation: 5

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is comprised of hard inplus, hard in, hard out, and hard exceed, as well as soft in and soft out. The offered counters cover traffic reclassified to in-profile (that includes traffic reclassified to inplus-profile), traffic reclassified to out-of-profile (that includes traffic reclassified to exceed-profile), and traffic that has not been reclassified at egress (uncolor). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

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 inplus**, **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 is 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 (hard in-profile) instead of offered-undefined (uncolored).

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in [Table 76](#) (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 76 Egress Accounting Statistics Collected in offered-profile-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-total-cir-exceed — Counter resource allocation: 3

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter. The **offered-total-cir-exceed** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-total-cir-exceed** mode is similar to the **offered-total-cir** mode except that it includes support for forwarded and dropped counters for **profile exceed**.

This mode is intended to be used when the policer is configured with **enable-exceed-pir** to forward packets that exceed its configured PIR or when traffic is egress reclassified to profile exceed. The mode gives the forwarded and dropped counters per profile (in, out, exceed). It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in [Table 77](#) (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 77 Egress Accounting Statistics Collected in offered-total-cir-exceed stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
Dro. ExcProf	xpd	ExceedProfilePktsDropped
	xod	ExceedProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded
For. ExcProf	xpf	ExceedProfilePktsForwarded
	xof	ExceedProfileOctetsForwarded

offered-four-profile-no-cir — Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering the policer. Offered, dropped, and forwarded counters are provided for inplus-profile, in-profile, out-of-profile, and exceed-profile traffic.

The **offered-four-profile-no-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-four-profile-no-cir** mode is similar to the **offered-profile-no-cir** mode except that it includes support for offered, dropped and forwarded counters for both profile inplus and profile exceed.

This mode is intended to be used when traffic is egress reclassified to profile inplus and/or exceed. It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in [Table 78](#) (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 78 Egress Accounting Statistics Collected in offered-four-profile-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. ExcProf	xpo	ExceedProfilePacketsOffered
	xoo	ExceedProfileOctetsOffered
Off. InplusProf	ppo	InplusProfilePacketsOffered
	poo	InplusProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
Dro. ExcProf	xpd	ExceedProfilePktsDropped
	xod	ExceedProfileOctetsDropped

Table 78 Egress Accounting Statistics Collected in offered-four-profile-no-cir stat-mode (Continued)

Show Output	Accounting Stats Collected	
	Field	Field Description
Dro. InplusProf	ppd	InplusProfilePktsDropped
	pod	InplusProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded
For. ExcProf	xpf	ExceedProfilePktsForwarded
	xof	ExceedProfileOctetsForwarded
For. InplusProf	ppf	InplusProfilePktsForwarded
	pof	InplusProfileOctetsForwarded

offered-total-cir-four-profile — Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. There is a separate dropped and forwarded counter for inplus, in, out, and exceed-profile traffic.

The **offered-total-cir-four-profile** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-total-cir-four-profile** mode is similar to the **offered-total-cir** except that it includes support for forwarded and dropped counters for both inplus-profile and exceed-profile.

This mode is intended to be used when traffic is egress reclassified to inplus-profile. It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in [Table 79](#) (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 79 Egress Accounting Statistics Collected in offered-total-cir-four-profile stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
Dro. ExcProf	xpd	ExceedProfilePktsDropped
	xod	ExceedProfileOctetsDropped
Dro. InprofProf	ppd	InplusProfilePktsDropped
	pod	InplusProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded
For. ExcProf	xpf	ExceedProfilePktsForwarded
	xof	ExceedProfileOctetsForwarded
For. InplusProf	ppf	InplusProfilePktsForwarded
	pof	InplusProfileOctetsForwarded

queue

Syntax `queue queue-id [queue-type] [create]`
`no queue queue-id`

Context `cfg>qos>qgrps>egr>queue-group`

Description	This command creates a queue for use in a queue group template. When created, the defined <i>queue-id</i> acts as a repository for the default parameters for the queue. The template queue is created on each queue group object that is created with the queue group template name. Each queue is identified within the template by a queue ID. 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.
Parameters	<i>queue-id</i> — Specifies the queue ID. The specified <i>queue-id</i> must exist within the egress queue group template. Values 1 to 8 <i>queue-type</i> — Specifies the method that the system uses to service the queue from a hardware perspective. Values expedite, best-effort 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. Default best-effort

7.12.2.2 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

queue-group

Syntax	queue-group [<i>queue-group-name</i>] { ingress egress } [{ association detail }] queue-group summary
Context	show>qos
Description	This command displays queue-group information.
Parameters	<i>queue-group-name</i> — Specifies the name of an existing queue group template up to 32 characters. ingress — Specifies whether the queue group name is an ingress policy. egress — Specifies whether the queue group name is an egress policy.

association — 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 instances per card (IOM or XCM).

Output The following outputs are examples of queue group information.

Sample Output

```
*A:PE# show qos queue-group egress
```

```
=====
Queue Group Egress
=====
Group-Name                Description
-----
qg1
_tmnx_nat_egr_q_grp        NAT/LNS Egress Queue Group Template
policer-output-queues     Default egress policer output queues.
_tmnx_nat_egr_q_grp_v2    NAT/LNS Egress Queue Group Template for ISAv2
_tmnx_lns_esm_egr_q_grp    LNS ESM Egress Queue Group Template
=====
Egress Queue Groups      : 5
=====
```

```
*A:PE#
*A:PE# show qos queue-group egress "qg1" detail
```

```
=====
QoS Queue-Group Egress
=====

-----
QoS Queue Group
-----
Group-Name      : qg1
Description     : (Not Specified)
Queues HQoS Managed : True
HS Attachment Policy : default
-----

Queue Information
-----
Queue-Id      : 1           Queue-Type      : best-effort
Admin PIR     : 50000       Admin CIR       : 0
PIR Rule      : closest     CIR Rule       : closest
CBS           : def         MBS              : 1000 KB
Packet Byte Offset : add 0
Named Buffer Pool : (Not Specified)
Adv Config Policy : (Not Specified)
Parent        : s1
PIR Level/Weight : 1/1       CIR Level/Weight : 0/1
Burst Limit    : default
Wred Queue Mode : none       Wred Queue SlopeUsage: not-applicable
Slope Policy   : default
Dynamic MBS    : enabled     Queue Delay     : not-applicable
```

```

High-Plus Drop Tail : def          High Drop Tail       : def
Low Drop Tail       : def          Exceed Drop Tail    : def

```

```

-----
HS Queue Information
-----

```

```

Queue-Id           : 1
Admin PIR           : 50000
PIR Rule            : closest
MBS                 : 1000 KB
Packet Byte Offset : add 0
Burst Limit         : default
HS Class Weight     : 1             HS Wrr Weight       : 1
HS Wred-Q Slope Plcy : _tmnx_hs_default
HS Alt Class Pool   : False

```

```

-----
HS Wrr Group Information
-----

```

```

HS Wrr Group Id    : 1             Class Weight       : 1
Rate               : max           Percent Rate        : 100.00
PIR Adaptation Rule: closest

HS Wrr Group Id    : 2             Class Weight       : 1
Rate               : max           Percent Rate        : 100.00
PIR Adaptation Rule: closest

```

```

=====
Queue Group FC Mapping
=====

```

```

FC Name                               Queue-Id
-----

```

```

No Matching Entries
=====

```

```

=====
Queue Group Ports (access)
=====

```

Port	Sched Pol	Acctg Pol	Stats	Description	QGrp-Instance
5/1/1		0	No		1

```

=====
Queue Group Ports (network)
=====

```

Port	Sched Pol	Policer-Ctrl-Pol	Acctg Pol	Stats	Description	QGrp-Instance
No Matching Entries						

```

=====
Qos Sap-Egress FC Group-Queue References
=====

```

```

Sap Policy      FC Name      Queue Id
-----
No Matching Entries

=====
Qos Sap-Egress FC Port-Redirect-Group-Queue References
=====
Sap Policy      FC Name      Queue Id
-----
No Matching Entries

=====
Queue Group Policer
=====
No Matching Entries

-----
HSMDB PIR Admin Packet WRR      MBS      Slope Plcy      WRR Plcy
Queue PIR Rule  Offset Weight      Max Class      Burst Lmt
-----
1      max      add 0      1      default      default      n/a
      closest      8      default
2      max      add 0      1      default      default      n/a
      closest      8      default
3      max      add 0      1      default      default      n/a
      closest      8      default
4      max      add 0      1      default      default      n/a
      closest      8      default
5      max      add 0      1      default      default      n/a
      closest      8      default
6      max      add 0      1      default      default      n/a
      closest      8      default
7      max      add 0      1      default      default      n/a
      closest      8      default
8      max      add 0      1      default      default      n/a
      closest      8      default
-----
=====
*A:PE#
*A:PE# show qos queue-group ingress

=====
Queue Group Ingress
=====
Group-Name      Description
-----
qg1
_tmnx_nat_ing_q_grp      NAT/LNS Ingress Queue Group Template
_tmnx_nat_ing_q_grp_v2    NAT/LNS Ingress Queue Group Template for ISA*
_tmnx_lns_esm_ing_q_grp    LNS ESM Ingress Queue Group Template
-----
Ingress Queue Groups      : 4
-----
=====
* indicates that the corresponding row element may have been truncated.
*A:PE#
*A:PE# show qos queue-group ingress "qg1" detail

=====
QoS Queue-Group Ingress

```

```

=====
-----
QoS Queue Group
-----
Group-Name      : qg1
Description     : (Not Specified)
-----
Q  Mode      CIR Admin  PIR Admin  CBS                      PIR Lvl/Wt Parent
              CIR Rule   PIR Rule   MBS                      CIR Lvl/Wt BurstLimit(B)
              Named-Buffer Pool      Pkt Bt Ofst      Adv Config Policy Name
              Low Drop Tail
              FIR Admin  FIR Rule   Cir-non-profiling
-----
No Matching Entries

=====
-----
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
-----
1              1              1              Access
-----
Entries found: 1
-----

=====
-----
Queue Group Service Maps
=====
Service Id              Service Type              Instance
-----
No Matching Entries

=====
-----
Queue Group Policer
=====
Policer Id      : 1
Description     : (Not Specified)
PIR Adptn       : closest                      CIR Adptn       : closest
Parent          : none                      Level           : 1
Weight          : 1                      Adv. Cfg Plcy   : none
Admin PIR       : max                      Admin CIR       : 0
CBS             : def                      MBS             : def
Hi Prio Only    : def                      Pkt Offset      : 0
Profile Capped  : Disabled

```

```

StatMode      : minimal

=====
*A:PE#
*A:PE# 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             1             0
2             0             1             0             0             0
3             0             0             0             0             0
4             0             0             0             0             0
5             0             3             0             0             0
-----
Total ingress QG templates per system : 4
Total egress  QG templates per system : 5
=====
*A:PE#

```

queue-group-redirect-list

- Syntax** `queue-group-redirect-list [policy-name] [association | detail]`
- Context** `show>qos`
- Description** This command displays queue group redirect list information.
- Parameters**
- policy-name* — Specifies the name of a queue group redirect list, up to 32 characters.
 - association** — Displays the entities associated with the specified queue group redirect list.
 - detail** — Displays detailed information, including queue group redirect list associations.
- Output** The following output is an example of queue group redirect list information.

Sample Output

```

A:PE# show qos queue-group-redirect-list detail
=====
Queue Group Redirect List Information
=====
Policy Name      : list1
Match Type       : vxlan-vni
-----
Match            Instance
-----
1                1
2                2
3                3
-----
-----

```

```

Service Associations
-----
Service ID      Service Type      SAP
-----
1                IES                1/1/1          (Ingress/Egress)
-----
=====
A:PE#

```

pools

Syntax **pools** *mda-id[/port]* [*access-app* [*pool-name* | **service** *service-id* | **queue-group** *queue-group-name* [**instance** *instance-id*]]]
pools *mda-id[/port]* [*network-app* [*pool-name* | **queue-group** *queue-group-name* [**instance** *instance-id*]]]
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 XMA or MDA.

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 to 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

instance *instance-id* — Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

Output The following output is an example of pool information.

Sample Output

```

*A:PE# show pools access-egress 5/1/1 queue-group "qg1" instance 1
=====
Pool Information
=====
Port                : 5/1/1
Application          : Acc-Egr           Pool Name           : default
CLI Config. Resv CBS : 30%(default)
Resv CBS Step        : 0%                Resv CBS Max         : 0%

```



```

Amber Alarm Threshold: 0%                      Red Alarm Threshold : 0%
-----
Queue-Groups
-----
Queue-Group:Instance : qgl:1
-----
Utilization          State      Start-Avg      Max-Avg      Max-Prob
-----
HiPlus-Slope         Down          85%           100%         80%
High-Slope           Down          70%           90%          80%
Low-Slope            Down          50%           75%          80%
Exceed-Slope         Down          30%           55%          80%
Time Avg Factor      : 7
Pool Total           : 75264 KB
Pool Shared          : 52224 KB      Pool Resv      : 23040 KB
-----
Current Resv CBS     Provisioned     Rising         Falling        Alarm
%age                 all Queues     Alarm Thd      Alarm Thd      Color
-----
30%                  0 KB          NA             NA             Green
Pool Total In Use    : 0 KB
Pool Shared In Use   : 0 KB      Pool Resv In Use : 0 KB
WA Shared In Use     : 0 KB
HiPlus-Slope Drop Pr*: 0      Hi-Slope Drop Prob : 0
Lo-Slope Drop Prob   : 0      Excd-Slope Drop Prob : 0
=====
Queue : accQGrp->qgl:1(5/1/1)->1
=====
FC Map               : not-applicable
Tap                  : not-applicable  FP              : not-applicable
Admin PIR            : 10000000    Oper PIR        : 50000
Admin CIR            : 0          Oper CIR        : 0
Admin MBS            : 122880 KB   Oper MBS        : 122880 KB
High-Plus Drop Tail  : 122880 KB   High Drop Tail  : 122880 KB
Low Drop Tail        : 110592 KB   Exceed Drop Tail : 98304 KB
CBS                  : 0 KB      Depth           : 0
Slope                : not-applicable
=====
* indicates that the corresponding row element may have been truncated.
*A:PE#

```

port

Syntax	port queue-group [{ ingress egress }] [<i>queue-group-name</i>] [{ access network }] [instance <i>instance-id</i>] [{ statistics associations queue-depth [<i>queue queue-id</i>]}] port queue-group summary
Context	show>port
Description	This command displays physical port information for the port's queue group.
Parameters	ingress — Specifies whether the queue group name is an ingress policy. 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.

access — Specifies that the queue group name is on an access port.

network — Specifies that the queue group name is on a network port.

instance-id — Specifies the identification of a specific instance of the queue group.

statistics — Displays statistical information for the queue group.

associations — Displays the entities associated with the specified queue group name.

queue-depth — Displays the queue depth information for the specified queue group queue.

queue-id — Specifies the identification of a queue.

Output The following output is an example of port information for the port's queue group.

Sample Output

```
*A:PE# show port 5/1/1 queue-group ingress
```

```
=====
Ethernet Port 5/1/1 Access Ingress Queue-Group
=====
Group Name      : qgl
Description     : (Not Specified)
Sched Policy    : None          Acct Pol      : None
Collect Stats   : disabled

Queues
-----
Ing. QGroup     : qgl          Queue-Id    : 1 (Unicast) (Priority)
Description     : (Not Specified)
Admin PIR       : max*         Admin CIR    : 0*
PIR Rule        : closest*     CIR Rule     : closest*
Admin FIR       : 0*           FIR Rule     : closest*
CBS             : def*         MBS          : def*
Monitor Depth   : not-applicable
Low Drop Tail   : def*
```

* means the value is inherited

```
=====
*A:PE#
```

```
*A:PE# show port 5/1/1 queue-group egress
```

```
=====
Ethernet port 5/1/1 Access Egress queue-group
=====
Group Name      : qgl          Instance-Id  : 1
Description     : (Not Specified)
Sched Policy    : None          Acct Pol     : None
Collect Stats   : disabled     Agg. Limit   : max
Limit Unused BW : Disabled
Frame Based Acc : Disabled
HS Turbo Queues : Disabled
```

```

Queues
-----
Queue-Group      : qg1           Instance-Id   : 1       Queue-Id    : 1
Description      : (Not Specified)
Admin PIR        : 50000*        Admin CIR     : 0*
PIR Rule         : closest*      CIR Rule      : closest*
CBS              : def*          MBS           : 1000 KB*
Weight           : not-applicable CIR Weight     : not-applicable
Monitor Depth    : not-applicable
HiPlus Drop Tail : def*          Hi Drop Tail  : def*
Low Drop Tail    : def*          Exceed Drop * : def*

* means the value is inherited
=====

Host-Matches
-----
No Host-Matches found
-----

Group Name       : policer-output-qu* Instance-Id   : 1
Description      : (Not Specified)
Sched Policy     : None          Acct Pol       : None
Collect Stats    : disabled      Agg. Limit     : max
Limit Unused BW  : Disabled
Frame Based Acc  : Disabled
HS Turbo Queues  : Disabled

Queues
-----
Queue-Group      : policer-output-qu* Instance-Id   : 1       Queue-Id    : 1
Description      : Default egress policer output queues.
Admin PIR        : max*          Admin CIR     : 0*
PIR Rule         : closest*      CIR Rule      : closest*
CBS              : def*          MBS           : def*
Weight           : not-applicable CIR Weight     : not-applicable
Monitor Depth    : not-applicable
HiPlus Drop Tail : def*          Hi Drop Tail  : def*
Low Drop Tail    : def*          Exceed Drop * : def*

Queue-Group      : policer-output-qu* Instance-Id   : 1       Queue-Id    : 2
Description      : Default egress policer output queues.
Admin PIR        : max*          Admin CIR     : 0*
PIR Rule         : closest*      CIR Rule      : closest*
CBS              : def*          MBS           : def*
Weight           : not-applicable CIR Weight     : not-applicable
Monitor Depth    : not-applicable
HiPlus Drop Tail : def*          Hi Drop Tail  : def*
Low Drop Tail    : def*          Exceed Drop * : def*

* means the value is inherited
=====

Host-Matches
-----
No Host-Matches found
-----
=====
* indicates that the corresponding row element may have been truncated.

```

```
*A:PE#
*A:PE# show port 5/1/1 queue-group summary

=====
Port queue-group summary
=====
Access-egress queue groups:
-----
qg1
policer-output-queues
Total number of access-egress queue groups : 2

Network-egress queue groups:
-----
Total number of network-egress queue groups : 0

Access-ingress queue groups:
-----
qg1
Total number of access-ingress queue groups : 1
=====
*A:PE#
```

7.12.2.3 Monitor Commands

card

Syntax	card <i>slot-number</i> fp <i>fp-number</i> ingress { access network } queue-group <i>queue-group-name</i> instance <i>instance-id</i> [interval <i>seconds</i>] [repeat <i>repeat</i>] policer <i>policer-id</i> [absolute percent-rate [<i>reference-rate</i>]]
Context	monitor
Description	This command monitors policer statistics in an ingress FP queue group.
Parameters	<i>slot-number</i> — Specifies the slot number associated with the queue group, expressed as an integer. Values 1 to 20 <i>fp-number</i> — Specifies the FP number associated with the queue group, expressed as an integer. Values 1 to 8 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. <i>queue-group-name</i> — Specifies the name of the queue group up to 32 characters.

instance-id — Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

seconds — Configures the interval for each display in seconds.

Default 11 seconds

Values 11 to 60

repeat — Configures how many times the command is repeated.

Default 10

Values 1 to 999

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

Values 1 to 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 <i>slot-number</i> fp <i>fp-number</i> queue-group <i>queue-group-name</i> instance <i>instance-id</i> [ingress] [access networks] [interval <i>seconds</i>] [repeat <i>repeat</i>] [absolute percent-rate [<i>reference-rate</i>]] [arbiter root <i>name</i>]
Context	monitor>qos>arbiter-stats
Description	This command monitors arbiter statistics in an ingress FP queue group.
Parameters	<i>slot-number</i> — Specifies the slot number associated with the queue group, expressed as an integer. Values 1 to 20 <i>fp-number</i> — Specifies the FP number associated with the queue group, expressed as an integer. Values 1 to 8 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-name* — Specifies the name of the queue group up to 32 characters.
- instance-id* — Specifies the identification of a specific instance of the queue-group.
- Values** 1 to 65535
- seconds* — Configures the interval for each display in seconds.
- Default** 11 seconds
- Values** 11 to 60
- repeat* — Configures how many times the command is repeated.
- Default** 10
- Values** 1 to 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
- 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, 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 <i>port-id</i> egress <i>network</i> queue-group <i>queue-group-name</i> instance <i>instance-id</i> [interval <i>seconds</i>] [repeat <i>repeat</i>] [absolute rate] [arbiter [root <i>name</i>]]
Context	monitor>qos>arbiter-stats
Description	This command monitors arbiter statistics in an egress port queue group.

Parameters	port	<i>port-id</i> —		
		<i>port-id</i>	<i>slot/mda/port [.channel]</i>	
		eth-sat-id	esat-id/slot/port	
			esat	keyword
			id	1 to 20
		pxc-id	pxc-id.sub-port	
			pxc	keyword
			id	1 to 64
			sub-port	a, b

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.

instance *instance-id* — Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

interval *seconds* — Configures the interval for each display in seconds.

Default 11 seconds

Values 11 to 60

repeat *repeat* — Configures how many times the command is repeated.

Default 10

Values 1 to 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, 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*]

port *port-id* **vport** *name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]

- Context** monitor>qos>scheduler-stats
- Description** This command enables port traffic monitoring.
- Parameters** **port** *port-id* — Specifies the port ID.

Values			
<i>port-id</i>	<i>slot/mda/port</i> [<i>.channel</i>]		
	eth-sat-id	esat-id/slot/port	
		esat	keyword
pxc-id		<i>id</i>	1 to 20
		pxc-id.sub-port	
		pxc	keyword
		<i>id</i>	1 to 64
		sub-port	a, b

- queue-group** *queue-group-name* — Specifies the name of the queue group up to 32 characters.
- instance** *instance-id* — Specifies the identification of a specific instance of the queue-group.
- Values** 1 to 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 to 60
- repeat** *repeat* — Configures how many times the command is repeated.
- Default** 10
- Values** 1 to 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.

8 Scheduler QoS Policies

8.1 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 that is used to place the scheduler within the hierarchy. Three tiers of virtual schedulers are supported. Root schedulers are often 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, policers, 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, policer, 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.

8.1.1 Egress Port-Based Schedulers

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 policers and queues with more bandwidth than available on an egress port, the actual bandwidth distribution to the policers and 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 policers and 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 [Port-Level Virtual Scheduler Bandwidth Allocation Based on Priority and CIR](#).

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, policer, 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, policers, and schedulers attached to level 1 for above-CIR bandwidth. The low-priority distribution for level 1 above-CIR is reserved for all orphaned policers, queues, and schedulers on the egress port. Orphans are policers, 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 policer or queue. The bandwidth allocated to the policer or queue is converted to a value for the PIR (maximum rate) setting of the policer or queue. This way, the hardware schedulers operating at the egress port level will only schedule bandwidth for all policers or 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 policer or queue is inactive or has a limited offered load that is below its fair share (fair share is based on the bandwidth allocation a policer or 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 policer or queue PIR was set to zero (or near zero), the policer or 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 policers and queues, the system calculates a Minimum Information Rate (MIR) for each policer and queue. To calculate each policer or queue MIR, the system determines what the Fair Information Rate (FIR) of the queue or policer would be if that policer or 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.

8.1.1.1 Service/Subscriber or Multiservice Site Egress Port Bandwidth Allocation

The port-based egress scheduler can be used to allocate bandwidth to each service or subscriber or multiservice site associated with the port. While egress policers and queues on the service can have a child association with a scheduler policy on the SAP or multiservice site, all policers and queues must vie for bandwidth from an egress port. Two methods are supported to allocate bandwidth to each service or subscriber or multiservice site queue:

1. Service or subscriber or multiservice site queue association with a scheduler on the SAP or multiservice site, either of which is associated with a port-level scheduler.
2. Service or subscriber or multiservice site queue association directly with a port-level scheduler.

Figure 19 Port-Level Virtual Scheduler Bandwidth Allocation Based on Priority and CIR




OSSG130

8.1.1.2 Service or Subscriber or Multiservice Site Scheduler Child to Port Scheduler Parent

The service or subscriber or multiservice site scheduler to port scheduler association model allows for multiple services or subscribers or multiservice sites 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 20](#) shows this model.

Figure 20 Two-Scheduler Policy Model for Access Ports




OSSG131

When 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 or multiservice site level schedulers and the policers and 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 or multiservice site 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, with the goal of a carrier to provide segments of bandwidth to providers who purchase that bandwidth as services. While the carrier does not care about the interior services of the provider, it does 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 or multiservice site context through the service or subscriber or multiservice site level policy, without affecting the overall bandwidth allocation to each service or subscriber or multiservice site. [Figure 21](#) 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 21 Schedulers on SAP or Multiservice Site Receive Bandwidth From Port Priority Levels



OSSG132


8.1.1.2.1 Direct Service or Subscriber or Multiservice Site 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 or multiservice site policer or 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 or multiservice site policer or 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 or multiservice site queuing. [Figure 22](#) shows the single, port scheduler policy model.

[Figure 22](#) also shows the optional aggregate rate limiter at the SAP, multiservice site or subscriber or multiservice site level. The aggregate rate limiter is used to define a maximum aggregate bandwidth at which the child queues and policers, if used, 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 or multiservice site is monitored. When the aggregate rate limit is reached, no more bandwidth is allocated to the children associated with the SAP, multiservice site, or subscriber or multiservice site. Aggregate rate limiting is restricted to the single scheduler policy model and is mutually exclusive to defining SAP, multiservice site, or subscriber or multiservice site 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 or multiservice site 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 22 Direct Service or Subscriber or Multiservice Site Association to Port Scheduler Model



OSSG133

8.1.1.3 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 or multiservice site. The service or subscriber or multiservice site 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 policer or queue accounted for. When considering HDLC encoded PoS or SDH ports on the 7750 SR, 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 policer or queue can use. This is based on the offered load the policer or queue is currently experiencing (how many octets are being offered). The offered load is compared to the configured CIR and PIR of the queue or policer. 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 policers or 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 ($\text{offeredPackets} \times 20 + \text{offeredOctets} = \text{frameOfferedLoad}$). This frame-offered-load value represents the amount of line rate bandwidth the policer or 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 on the 7750 SR, 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 ($\text{offeredOctets} \times \text{configuredEncapsulationOverhead} + \text{offeredOctets} = \text{frameOfferedLoad}$). The frame-offered-load value is used by the egress PoS/SDH port scheduler in the same manner as the egress Ethernet port scheduler.

From a provisioning perspective, queues, policers, 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 23](#) and [Figure 24](#) show 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 23 Port Bandwidth Distribution for Service and Port Scheduler Hierarchies



OSSG134

Figure 24 Port Bandwidth Distribution for Direct Queue to Port Scheduler Hierarchy



OSSG135

8.1.1.4 Parental Association Scope

A **port-parent** command in the sap-egress and network-queue QoS policy policer or queue context defines the direct child/parent association between an egress policer or queue and a port scheduler priority level. The **port-parent** command is mutually exclusive to the already-existing **scheduler-parent** or **parent** command, which associates a policer or queue with a scheduler at the SAP, multiservice site, or subscriber or multiservice site profile level. It is possible to mix local parented (parent to service or subscriber or multiservice site level scheduler) and port parented policers and 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 policer or queue with a defined parent port is on a port without a port scheduler policy applied, that policer or queue will be considered orphaned. If a policer or queue with a **scheduler-parent** or **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 policer or queue will be considered orphaned as well.

A queue or policer can be moved from a local parent (on the SAP, multiservice site, or subscriber or multiservice site profile) to a port parent priority level simply by executing the **port-parent** command. When the **port-parent** command is executed, any local parent information for the policer or queue is lost. The policer or queue can also be moved back to a local scheduler-parent or parent at any time by executing the **scheduler-parent** or **parent** command. Lastly, the local scheduler parent, parent, or port parent association can be removed at any time by using the **no** form of the appropriate parent command.

8.1.1.5 Service or Subscriber or Multiservice Site-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, the **scheduler-parent** command, and the **parent** command for schedulers at tiers 2 and 3 within the scheduler policy are mutually exclusive. 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. When the **port-parent** command is executed, any local scheduler-parent or parent information for the scheduler is lost. The schedulers at tiers 2 and 3 can also be moved back to a local scheduler-parent or parent at any time by executing the local **scheduler-parent** or **parent** command. Lastly, the local scheduler parent, parent, or port parent association can be removed at any time by using the **no** form of the appropriate parent command. A scheduler in tier 1 with a port parent definition can be added or removed at any time.

8.1.1.6 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 25](#) shows port-based virtual scheduling bandwidth distribution.

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

Figure 25 Bandwidth Distribution on Network Port with Port-Based Scheduling

8.1.1.7 Foster Parent Behavior for Orphaned Queues and Schedulers

All queues, policers, 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, policers, 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, multiservice site, subscriber and port level.

By default, orphaned queues, policers, and schedulers are allocated bandwidth after all queues, policers, and schedulers in the parented hierarchy have had bandwidth allocated within-CIR and above-CIR. Therefore, an orphaned scheduler, policer, or queue can be considered as being foster parented by the port scheduler. Orphaned queues, policers, 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 policers, 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 when all other properly parented queues, policers, and schedulers have received bandwidth. Orphaned queues, policers, 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, policers, and schedulers are treated in a similar fashion as other properly parented queues, policers, and schedulers based on the override parenting parameters.


It is expected that an orphan condition is not the desired state for a queue, policer, or scheduler and is the result of a temporary configuration change or configuration error.

8.1.1.8 Congestion Monitoring on Egress Port Scheduler

A typical example of congestion monitoring on an Egress Port Scheduler (EPS) is when the EPS is configured within a Vport. A Vport is a construct in an HQoS hierarchy that can be used to control the bandwidth associated with an access network element (such as, GPON port, OLT, DSLAM) or a retailer that has subscribers on an access node (among other retailers).

The example in [Figure 26](#) shows Vports representing GPON ports on an OLT. For capacity planning purposes, it is necessary to know if the GPON ports (Vports) are congested. Frequent and prolonged congestion on the Vport will prompt the operator to increase the offered bandwidth to its subscribers by allocating additional GPON ports and subsequently moving the subscribers to the newly allocated GPON ports.

Figure 26 GPON Bandwidth Control through Vport



al_0740

There are no forward/drop counters directly associated with the EPS. Instead, the counters are maintained on a per queue level. Consequently, any indication of the congestion level on the EPS is derived from the queue counters that are associated with the given EPS.


The EPS congestion monitoring capabilities rely on a counter that records the number of times that the offered EPS load (measured at the queue level) crossed the predefined bandwidth threshold levels within a given, operator defined timeframe. This counter is called the exceed counter. The rate comparison calculation (offered rate vs threshold) are executed several times per second and the calculation interval cannot be influenced externally by the operator.

The monitoring threshold can be configured via CLI per aggregate EPS rate, EPS level or EPS group. The threshold is applicable to PIR rates.

To enable congestion monitoring on EPS, monitoring must be explicitly enabled under the Vport object itself or under the physical port when the EPS is attached directly to the physical port. In addition, the monitoring threshold within the EPS must be configured.


Two examples of congestion monitoring on an EPS that is configured under the Vport are shown in [Figure 27](#) and [Figure 28](#). [Figure 28](#) shows more severe congestion than [Figure 27](#). The EPS exceed counter (the number of dots above the threshold line) can be obtained via a CLI show command or read directly via MIBs.

Figure 27 Exceed Counts



al_0741

Figure 28 Exceed Counts (Severe Congestion)




al_0742

When the exceed counter value is obtained, the counter should be cleared, which resets the exceed counter and number of samples to zero. This is because the longer the interval between a clear and a show or read, the more diluted the congestion information becomes. For example, 100 threshold exceeds within a 5 minute interval depicts a more accurate congestion picture compared to 100 threshold exceeds within a 5 hour interval.


The reduced ability to determine the time of congestion if the reading interval is too long is shown in [Figure 29](#), [Figure 30](#), and [Figure 31](#). It can be seen that the same readings (in the 3 examples) can represent different congestion patterns that occur at different times between the two consecutive reads. The congestion pattern, or the exact time of congestion cannot be determined from the reading itself. The reading only indicates that the congestion occurred x number of times between the two consecutive readings. In the example shown in [Figure 29](#), [Figure 30](#), and [Figure 31](#), an operator can decipher that the link was congested 20% of the time during a one day period without being able to pinpoint the exact time of congestion within the one day period. To determine the time of the congestion more accurately, the operator must collect the information more frequently. For example, if the information is collected every 30 minutes, then the operator can determine the part of the day during which congestion occurred within 30 minutes of accuracy.

Figure 29 **Determining the Time of Congestion (Example 1)**




al_0743

Figure 30 **Determining the Time of Congestion (Example 2)**



al_0744

Figure 31 Determining the Time of Congestion (Example 3)



al_0748

8.1.1.8.1 Scalability, Performance, and Operation

The scalability and performance is driven by the number of entities for which congestion monitoring is enabled on each line card.

Each statistics gathering operation requires a show or read followed by a clear. The shorter the time between the two, the more accurate the information about the congestion state of the EPS will be.

If the clear operation is not executed after the show or read operation, the external statistics gathering entity (external server) would need to perform additional operations (such as, subtract statistics between the two consecutive reads) in order to obtain the delta between the two reads.

The recommended minimum polling interval at a higher scale (high number of monitoring entities) is 15 minutes per monitoring entity.

If statistics are obtained via SNMP, the relevant MIB entries corresponding to the show command are:

- tPortEgrVPortMonThrEntry
- tPortEgrMonThrEntry

Clearing of the statistics can also be performed through a common MIB entry, corresponding to a clear command: tmnxClearEntry.

8.1.1.8.2 Restrictions

The scalability and performance is driven by the number of entities for which congestion monitoring is enabled on each line card.

- EPS congestion monitoring is supported only on Ethernet ports.
- If EPS is applied under Vports, the congestion monitoring mechanism does not provide any indication of whether or not the physical port was congested. For example, if the physical port is congested, it will distribute less bandwidth to its Vports than it otherwise would. Therefore, the Vport offered load will appear as less than it actually is, giving an impression of no congestion.
- Changing EPS parameters dynamically does not automatically update congestion monitoring statistics. Therefore, it is recommended that the congestion monitoring statistics are cleared after changing the values of EPS parameters.
- If EPS is configured under a LAG, the failure of the active link causes an interruption in statistics collection.

8.1.2 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 that 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 or policer 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 policer or queue operates in frame-based accounting mode, and must be managed by a virtual scheduler policy or by a port virtual scheduler policy. Egress policers or 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 policer or queue is accomplished through a frame-based accounting command defined on the scheduling policy level associated with the policer or queue or through a policer or queue frame-based accounting parameter on the aggregate rate limit command associated with the queues SAP, multiservice site or subscriber or multiservice site context. Packet byte offset settings are not included in the applied rate when frame-based accounting is configured, however the offsets are applied to the statistics.

8.1.2.1 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 defined CIR and PIR threshold rates of the policer or queue to determine bandwidth offerings from the policer or queue virtual scheduler. Frame-based usage on the wire allows the port bandwidth to be accurately allocated to each child policer and queue on the port.

8.1.2.2 Existing Egress Port-Based Virtual Scheduling

The port-based virtual scheduling mechanism takes the native packet-based accounting results from the policer or queue and adds 20 bytes to each packet to derive the frame-based offered load of the policer or queue. 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 policer or queue. When the port virtual scheduler computes the amount of bandwidth allowed to the policer or queue (in a frame-based fashion), the bandwidth is converted back to a packet-based value and used as the operational PIR or the policer or queue. The native packet-based mechanisms of the policer or queue continue to function, but the maximum operational rate is governed by frame-based decisions.

8.1.2.3 Behavior Modifications for Frame-Based Accounting

The frame-based accounting feature extends this capability to allow the policer or queue CIR and PIR thresholds to be defined as frame-based values as opposed to packet-based values. The policer or 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 operational packet-based CIR and PIR of the policer or queue are accurately modified during each iteration of the virtual scheduler to represent the provisioned frame-based CIR and PIR. Packet byte offset settings are not included in the applied rate when frame-based accounting is configured, however the offsets are applied to the statistics.

8.1.2.4 Virtual Scheduler Rate and Queue Rate Parameter Interpretation

Normally, a scheduler policy contains rates that indicate packet-based accounting values. When the children 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 or policer frame-based operation, enabling frame-based operation is configured at either the scheduling policy or aggregate rate limit command level. All policers and 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 policer and queue 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 100 Mb/s on two different SAPs, one associated with a policy with frame-based accounting enabled and the other without frame-based accounting enabled, the 100 Mb/s 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. Packet byte offset settings are not included in the applied rate when frame-based accounting is configured; however, the offsets are applied to the statistics.

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 policers and queues managed by the aggregate rate limit and converts them from packet-based to frame-based accounting mode.

8.1.3 Virtual Scheduling Unused Bandwidth Distribution

The Hierarchical QoS (HQoS) mechanism is designed to enforce a user definable hierarchical shaping behavior on an arbitrary set of policers and queues. The mechanism accomplishes this by monitoring the offered rate of each policer and queue and using the result as an input to a virtual scheduler hierarchy defined by the user. The hierarchy consists of a number of virtual scheduler with configurable maximum rates per scheduler and attachment parameters between each. The parameters consist of weights and priority levels used to distribute the available bandwidth in a top-down fashion through the hierarchy with the queues at the bottom. The resulting bandwidth provided to each member policer and queue by the virtual schedulers is then configured as an operational PIR on the corresponding hardware policer or queue, which prevents that policer or queue from receiving more hardware scheduler bandwidth than dictated by the virtual scheduler.

8.1.3.1 Default Unused Bandwidth Distribution

The default behavior of HQoS is to only throttle active policers and queues currently exceeding their allocated bandwidth. by the virtual schedulers controlling the active policer or queue. A policer or queue that is currently operating below its share of bandwidth is allowed an operational PIR greater than its current rate; this includes inactive policers and queues. The operational PIR for a policer or queue is capped by its admin PIR and set to the fair-share of the available bandwidth for the policer or queue based on its priority level in the HQoS hierarchy and its weight within that priority level. The result is that between HQoS iterations, a policer or queue below its share of bandwidth may burst to a higher rate and momentarily overrun the prescribed aggregate rate.

This default behavior works well in situations where an aggregate rate is being applied as a customer capping function to limit excessive use of network resources. However, in certain circumstances where an aggregate rate must be maintained due to limited downstream QoS abilities or due to downstream priority-unaware aggregate policing, a more conservative behavior is required. The following functions can be used to control the unused bandwidth distribution:

- The **above-offered-cap** command within the **adv-config-policy** provides control of the operational PIR to prevent aggregate rate overrun of each policer or queue. This is accomplished by defining how much the operational PIR of a queue or policer is allowed to exceed its current allocated bandwidth.
- The **limit-unused-bandwidth** (LUB) command.

8.1.3.2 Limit Unused Bandwidth

The limit-unused-bandwidth (LUB) command protects against exceeding the aggregated bandwidth by adding a LUB second-pass to the HQoS function, which ensures that the aggregate fair-share bandwidth does not exceed the aggregate rate.

The command can be applied on any tier 1 scheduler within an egress scheduler policy or within any agg-rate node (except when using the HSMDB) and affects all policers and queues controlled by the object.

When LUB is enabled, the LUB second pass is performed as part of the HQoS algorithm. The order of operation between HQoS and LUB is as follows:

- Policer or queue offered rate calculation.
- Offered rate modifications based on **adv-config-policy offered-measurement** parameters.
- HQoS Bandwidth determination based on modified offered-rates.
- LUB second pass to ensure aggregate rates are not exceeded where LUB enabled.
- Bandwidth distribution modification based on **adv-config-policy bandwidth-distribution** parameters.
- Each policer or queue operational PIR is then modified.

When LUB is enabled on a scheduler rate or aggregate rate, a LUB context is created containing the rate and the associated policers and queues that the rate controls. Because a policer or queue may be controlled by multiple LUB enabled rates in a hierarchy, the policer or queue may be associated with multiple LUB contexts.

LUB is applied to the contexts where it is enabled. LUB first considers how much of the aggregate rate is unused by the aggregate rates of each member policer or queue after the first pass of the HQoS algorithm. This represents the current bandwidth that may be distributed between the member. LUB then distributes the available bandwidth to its members based on each LUB-weight. The LUB-weight is determined as follows:

- If a policer or queue is using all of its default HQoS assigned rate, then its LUB-weight is 0. It is not participating in the bandwidth distribution as it cannot accept more bandwidth.


- If a policer or queue has accumulated work, its LUB-weight is set to 50. The work is determined by the queue having built up a depth of packets, or the offered rate of a policer or queue increasing since last sample period. The aim is to assign more of the unused bandwidth to policers and queues needing more capacity. No attempt is made to distribute based on the relative priority level or weight within the hierarchy of a policer or queue.
- Otherwise, a queue's LUB-weight is 1

The resulting operational PIRs are then set such that the scheduler or agg-rate rate is not exceeded. To achieve the best precision, policers and queues must be configured to use **adaptation-rule pir max cir max** to prevent the actual rate used exceeding that determined by LUB.

Example

For a simple scenario with 5 egress SAP queues all without rates configured but with each queue parented to a different level in a parent scheduler that has a rate of 100Mb/s, see [Figure 32](#).

Figure 32 Limit Unused Bandwidth Example



The resulting bandwidth distribution is shown in [Figure 33](#); first, when no traffic is being sent with and without LUB applied, then when 20 Mb/s and 40 Mb/s are sent on queues 3 and 5, respectively, again with and without LUB applied. As shown, the distribution of bandwidth in the case where traffic is sent and LUB is enabled is based on the LUB-weights.

Figure 33 Resulting Bandwidth Distribution

No Traffic			Traffic in Q5 & Q3		
Offered Traffic	Default H-Qos	Second Pass (LUB)	Offered Traffic	Default H-Qos	Second Pass (LUB)
0	100	20	40	100	$59.39=40+19.39$
0	100	20	0	60	$0.39=0+0.39$
0	100	20	20	60	$39.42=20+19.42$
0	100	20	0	40	$0.39=0+0.39$
0	100	20	0	40	$0.39=0+0.39$

Each queue can burst to the full available capacity.

Each queue can burst to 1/5th of the available capacity.
[100/5=20]

Each queue can burst to the available capacity at that level. (full capacity minus capacity used by higher levels)

Each queue can burst to the used capacity plus its allocated part of the unused capacity.

al_0611

8.1.4 Configuring Port Scheduler Policies

8.1.4.1 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, policers, 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 or policer offered load, within their defined CIR. An overall maximum rate parameter defines the total bandwidth that will be allocated to all priority levels.

8.1.4.2 Special Orphan Queue and Scheduler Behavior

When a port scheduler is present on an egress port or channel, the system ensures that all policers, queues, and schedulers receive bandwidth from that scheduler to prevent free-running policers or queues that can cause the aggregate operational PIR of the port or channel to oversubscribe the bandwidth available. When the aggregate maximum rate for the policers and queues on a port or channel operates

above the available line rate, the forwarding ratio between the policers and queues will be affected by the hardware schedulers on the port and may not reflect the scheduling defined on the port or intermediate schedulers. Policers, 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, policers, 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 and, by default, are tied to priority 1 with a weight of 0. All weight 0 policers, 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.

8.1.4.3 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 policers, queues, and intermediate service-level schedulers (a service-level scheduler is a scheduler instance at the SAP, multiservice site, or subscriber or multiservice site profile level) operate based on packet overhead that 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 policer, 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 or policer 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 policer or queue to frame-CIR and frame-PIR values. When a service-level scheduler hierarchy is between the policers, 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 policer and queue level.

Packet over SONET (PoS) and SDH queues on the 7450 ESS and 7750 SR 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.

8.1.4.4 Aggregate Rate Limits for Directly Attached Queues

When all policers and queues for a SAP, multiservice site or subscriber or multiservice site 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 or multiservice site and the agg-rate limit provides this ability. Queues and policers 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 policers or 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 or multiservice site profile level. When an aggregate rate limit is in effect, a scheduler policy cannot be assigned. When a scheduler policy is assigned on the egress side of a SAP or subscriber or multiservice site profile, an agg-rate limit cannot be assigned.

Since the sap-egress policy defines a policer or queue parent association before the policy is associated with a service SAP or subscriber or multiservice site 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. Policers and queues in this state are considered to be orphaned and automatically attached to port scheduler priority 1. Orphaned policers and queues are included in the aggregate rate limiting behavior on the SAP or subscriber or multiservice site instance they are created within.

8.1.4.5 SAP Egress QoS Policy Queue Parenting

A SAP-egress QoS policy policer or queue may be associated with either a port parent or an intermediate scheduler parent. The validity of the parent definition cannot be checked at the time that it is provisioned since the application of the QoS policy is not known until it is applied to an egress SAP, subscriber, or multiservice site profile. Port or intermediate parenting can be decided on a queue-by-queue or policer-by-policer basis, some policers and queues tied directly to the port scheduler priorities while others are attached to intermediate schedulers.

8.1.4.6 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.

8.1.4.7 Egress Port Scheduler Overrides

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 1) has no configuration parameters and cannot be overridden.

8.1.4.8 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 (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.

This feature applies to the 7450 ESS and 7750 SR only.

[Figure 34](#) shows 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 34 Applying a Port Scheduler Policy to a Vport

The user adds a Vport to an Ethernet port using the following command:

CLI Syntax: `config>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. However, it 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: `config>port>ethernet>access>egress>vport>port-scheduler-policy port-scheduler-policy-name`

A Vport cannot be parented to the port scheduler when the Vport is using a port scheduler policy. It is important that 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 that uses an **agg-rate**, or a scheduler-policy, can be parented to a port scheduler. This is

explained in [Applying Aggregate Rate Limit to a Vport](#). The application of the **agg-rate rate**, **port-scheduler-policy**, and **scheduler-policy** commands under a Vport are mutually exclusive.

Each subscriber host policer or queue is port parented to the Vport that corresponds to the destination DSLAM using the existing port-parent command:

CLI Syntax:

```
config>qos>sap-egress>queue>port-parent [weight weight]
[level level] [cir-weight cir-weight] [cir-level cir-
level]
config>qos>sap-egress>policer>port-parent [weight
weight] [level level] [cir-weight cir-weight] [cir-
level cir-level]
```

This command can parent the policer or queue to either a port or to a Vport. These operations are mutually exclusive in CLI. When parenting to a Vport, the parent Vport for a subscriber host policer or queue is not explicitly indicated in the command; it is determined indirectly. The determination of the parent Vport for a given subscriber host queue is described in the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide*.

Subscriber host policers or queues, SLA profile schedulers, subscriber profile schedulers, and PW SAPs (in IES or VPRN services) can be parented to a Vport.

8.1.4.9 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 oversubscribe 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 *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide*.

8.1.4.10 Applying a Scheduler Policy to a Vport

The user can apply a scheduler policy to the Vport. This 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. This feature applies only to the 7450 ESS and 7750 SR.

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.

8.1.4.11 Weighted Scheduler Group in a Port Scheduler Policy

The port scheduler policy defines a set of eight priority levels. 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 group object is defined under the port scheduler policy, as follows:

CLI Syntax: `config>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 for the 7450 ESS and 7750 SR shown in the Vport in [Figure 34](#) consists of level priority 3 and level priority 4. Therefore, it inherits priority 4 when competing for bandwidth with the standalone priority levels 8, 7, and 5.

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: `config>qos>port-scheduler-policy>level priority-level
rate pir-rate [cir cir-rate] group group-name [weight
weight-in-group]`

The CLI will enforce that mapping of levels to a group are contiguous. In other words, a user would not be able to add a 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.

8.2 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.

8.2.1 Creating 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 1, 2, or 3 tier.
- Optionally, 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
ed
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#

8.2.2 Applying Scheduler Policies

Apply scheduler policies to the following entities:

- [Customer](#)
- [Epipe](#)
- [IES](#)
- [VPLS](#)
- [VPRN](#)

8.2.2.1 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
```

8.2.2.2 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

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#

```

8.2.2.3 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#

```

8.2.2.4 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#

```

8.2.2.5 VPRN

Use the following CLI syntax to apply scheduler policies to ingress and/or egress VPRN SAPs on the 7750 SR and 7950 XRS:

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-tunnel
        resolution-filter
        resolution-filter ldp
    vrf-target target:10001:1
    interface "to-cel" create
        address 192.168.0.0/24
        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#
```

8.2.3 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.

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#
```

8.2.4 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 the following contexts:

- **sap-egress>queue** *queue-id*
- **sap-egress>policer** *policer-id*
- **cfg>qos>qgrps>egr>qgrp>queue** *queue-id*
- **network-queue>queue** *queue-id*
- **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 policer, queue, or scheduler. If the port-parent command is executed without any parameters, the default parameters are assumed.

8.2.4.1 Within-CIR Priority Level Parameters

The within-CIR parameters define which port priority level the policer, queue, or scheduler should be associated with when receiving bandwidth for the policer, queue, or schedulers within-CIR offered load. The within-CIR offered load is the amount of bandwidth the policer, queue, or scheduler 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** command **cir-level** and **cir-weight** keywords. The **cir-level** keyword defines the port priority level that the scheduler, policer, or queue uses to receive bandwidth for its within-CIR offered load. The **cir-weight** is used when multiple queues, policers, or schedulers exist at the same port priority level for within-CIR bandwidth. The weight 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, policer, 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.

8.2.4.2 Above-CIR Priority Level Parameters

The above-CIR parameters define which port priority level the policer, queue, or scheduler should be associated with when receiving bandwidth for the above-CIR offered load of the policer, queue, or scheduler. The above-CIR offered load is the amount of bandwidth the queue, policer, or scheduler could use that is equal to or less than its defined PIR value (based on the queue, policer, or scheduler **rate** command) less any bandwidth that was given to the queue, policer, 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 policer, scheduler, or queue uses to receive bandwidth for its above-CIR offered load. The weight is used when multiple policers, queues, or schedulers exist at the same port priority level for above-CIR bandwidth. The weight 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]
policer policer-id [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]
```

8.2.5 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

```

In this sample, before enabling **dist-lag-rate-shared**, in the **port-scheduler-policy psp**, the max-rate achieved is twice 413202 kb/s (826 Mb/s). 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 kb/s. This can be confirmed using the following show output.

When `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, 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     yes       1         -      32768
3/2/1        up     active  up     yes       1         -      32768
3/2/2        up     active  up     yes       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 a percentage, 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 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

8.3 Service Management Tasks

This section discusses QoS scheduler policy service management tasks:

8.3.1 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 multiservice customer site causes the created schedulers to be removed, which makes them unavailable for SAP policers or queues associated with the customer site. Queues or policers that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler.

A QoS scheduler policy cannot be deleted until it is removed from all customer multiservice 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#
```

8.3.1.1 Removing a QoS Policy from a Customer Multiservice 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
```

8.3.1.2 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

```

The following syntax and example apply to the 7750 SR and 7950 XRS.

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
config>service>vprn# interface "to-cel"
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#

```

8.3.1.3 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`

8.3.2 Copying and Overwriting Scheduler Policies

An existing QoS policy can be copied, renamed with a new QoS policy value, or used to 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>gos# 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 servic
ed
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 servic
ed
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#

```

8.3.3 Editing QoS Policies

Existing policies and entries in the CLI can be edited. The changes are applied immediately to all customer multiservice 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, then overwrite the original policy.

8.4 QoS Scheduler Policy Command Reference

8.4.1 Command Hierarchies

- [Scheduler Policy Configuration Commands](#)
- [Port Scheduler Policy Configuration Commands](#)
- [Operational Commands](#)
- [Show Commands](#)
- [Clear Commands](#)

8.4.1.1 Scheduler Policy Configuration Commands

```

config
  — qos
    — scheduler-policy scheduler-policy-name [create]
    — no scheduler-policy scheduler-policy-name
      — description description-string
      — no description
      — [no] frame-based-accounting
      — [no] tier tier
        — parent-location {none | sub | vport}
        — no parent-location
        — scheduler scheduler-name [create]
        — 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

```

8.4.1.2 Port Scheduler Policy Configuration Commands

```

config
  — qos
    — port-scheduler-policy port-scheduler-name [create]

```

- **no port-scheduler-policy** *port-scheduler-name*
 - **description** *description-string*
 - **no description**
 - **[no] dist-lag-rate-shared**
 - **group** *name* [**create**]
 - **no group** *name*
 - **monitor-threshold** *percent*
 - **no monitor-threshold**
 - **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*] [**monitor-threshold** *percent*]
- **level** *priority-level* **percent-rate** *pir-percent* [**percent-cir** *cir-percent*] **group** *name* [**weight** *weight*] [**monitor-threshold** *percent*]
- **level** *priority-level* **rate** *pir-rate* [**cir** *cir-rate*] [**monitor-threshold** *percent*]
- **level** *priority-level* **percent-rate** *pir-percent* [**percent-cir** *cir-percent*] [**monitor-threshold** *percent*]
- **no level** *priority-level*
- **max-rate** *pir-rate*
- **max-rate** *percent* *percent-rate*
- **no max-rate**
- **monitor-threshold** *percent*
- **no monitor-threshold**
- **orphan-override** [**level** *level*] [**weight** *weight*] [**cir-level** *cir-level*] [**cir-weight** *cir-weight*]
- **no orphan-override**

8.4.1.3 Operational Commands

- ```

config
 — qos
 — copy scheduler-policy src-name dst-name [overwrite]
 — copy port-scheduler-policy src-name dst-name [overwrite]

```

### 8.4.1.4 Show Commands

- ```

show
  — qos
    — scheduler-hierarchy
      — customer customer-id site customer-site-name [scheduler scheduler-name] [ingress | egress] [detail]
      — port port-id queue-group queue-group-name [scheduler scheduler-name] [ingress | egress] [access | network] [instance instance-id] [detail]

```

```

— port port-id vport name summary
— port port-id [detail]
— port port-id vport name [scheduler scheduler-name] [detail]
— sap sap-id [scheduler scheduler-name] [ingress | egress] [detail]
— sap sap-id encap-group group-name [scheduler scheduler-name] [member
encap-id] [detail]
— subscriber sub-indent-string [scheduler scheduler-name] [ingress | egress]
[detail]
— subscriber sub-indent-string sla-profile sla-profile-name sap sap-id
[scheduler scheduler-name] [detail] [spi-sharing-type spi-sharing-type]
[spi-sharing-id spi-sharing-id]
— scheduler-name scheduler-name
— scheduler-policy [scheduler-policy-name] [association | sap-ingress policy-id | sap-
egress policy-id]
— scheduler-policy scheduler-policy-name {ingress | egress} queue-group queue-
grp-name
— scheduler-stats
— customer customer-id site customer-site-name [scheduler scheduler-name]
[ingress | egress]
— sap sap-id [scheduler scheduler-name] [ingress | egress]
— sap sap-id encap-group group-name [scheduler scheduler-name] [member
encap-id]
— subscriber sub-indent-string sap sap-id sla-profile sla-profile-name
[scheduler scheduler-name] [spi-sharing-type spi-sharing-type] [spi-
sharing-id spi-sharing-id]
— subscriber sub-indent-string [scheduler scheduler-name] [ingress | egress]

show
— qos
— agg-rate
— customer customer-id site customer-site-name [egress] [detail]
— port port-id queue-group queue-group-name [egress] [access | network]
[instance instance-id] [detail]
— port port-id vport name [detail]
— sap sap-id [egress] [detail]
— sap sap-id encap-group group-name [member encap-id] [detail]
— subscriber sub-indent-string [egress] [detail]

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

```

8.4.1.5 Clear Commands

```
clear
  — qos
    — scheduler-stats
      — sap sap-id [scheduler scheduler-name] [ingress | egress]
      — sap sap-id [scheduler scheduler-name] encap-group group-name [member
        encap-id]
```

8.4.2 Command Descriptions

8.4.2.1 Configuration Commands

8.4.2.1.1 Generic Commands

description

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

8.4.2.2 Scheduler Policy Commands

scheduler-policy

Syntax	scheduler-policy <i>scheduler-policy-name</i> [create] no scheduler-policy <i>scheduler-policy-name</i>
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 editing of an existing policy. The policy defines the hierarchy and operating parameters for virtual schedulers. 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 multiservice 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 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.

Parameters	<i>scheduler-policy-name</i> — The name of the scheduler policy.
Values	Valid names consist of any string up to 32 characters, 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	[no] frame-based-accounting

Context	config>qos>scheduler-policy
Description	The frame-based-accounting command is used to enable frame-based accounting for both the children queues parented to the scheduling policy and for the schedulers within the scheduler policy. When 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. Packet byte offset settings are not included in the applied rate when frame-based accounting is configured; however the offsets are applied to the statistics. The no form of this 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.

tier

Syntax	[no] tier <i>tier</i>
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	<i>tier</i> — 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 to 3

parent-location

Syntax	parent-location {none sub vport} no parent-location
Context	config>qos>scheduler-policy>tier
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 this command reverts to the default.
Default	parent-location 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 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.

scheduler

Syntax	scheduler <i>scheduler-name</i> [create] no scheduler <i>scheduler-name</i>
Context	config>qos>scheduler-policy>tier
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 which takes bandwidth from a scheduler in a higher tier. 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 <i>scheduler-name</i> 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 SLAs. If the <i>scheduler-name</i> 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 <i>scheduler-name</i> 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 <i>scheduler-name</i> 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, 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	<i>scheduler-name</i> — The name of the scheduler. Values Valid names consist of any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

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 aggregate 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>scheduler
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, policers (at egress only), 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 any time 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 this 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. When 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	<i>scheduler-name</i> — The <i>scheduler-name</i> must already exist within the context of the scheduler policy in a tier that is higher (numerically lower).
Values	Any valid <i>scheduler-name</i> existing on a higher tier within the scheduler policy.

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 zero (0) 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 in relation to their relative weights.

Values 1 to 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 zero (0) **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 to 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 0 to 8

Default 0

port-parent

Syntax	port-parent [weight <i>weight</i>] [level <i>level</i>] [cir-weight <i>cir-weight</i>] [cir-level <i>cir-level</i>] no port-parent
Context	config>qos>scheduler-policy>tier>scheduler
Description	The port-parent command defines a child/parent association between an egress scheduler and a port-based scheduler, or between an intermediate service scheduler and a port-based scheduler. 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 scheduler. If the port-parent command is executed without any parameters, the default parameters are assumed. In this context, the port-parent command and the parent command (used to create a parent/child association to an intermediate scheduler) are mutually exclusive. 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 policer or queue may cause a SAP or subscriber context of the policer or queue (policy associated with a SAP or subscriber profile) to enter an orphaned state. If an instance of a policer or queue is created on a port or channel that does not have a port scheduler enabled and the sap-egress policy creating the policer queue has a port-parent association, the policer or 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 policer or queue is on a port or channel that has a port scheduler configured and the sap-egress policy defines the policer or queue as having a non-existent intermediate scheduler parent, the policer or 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 scheduler. If a port scheduler is defined on the port that the scheduler instance exists, the 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 <i>weight</i> — Defines the weight the queue or scheduler will use at the above-CIR port priority level (defined by the level parameter). Values 0 to 100 Default 1 level <i>level</i> — Defines the port priority the queue or scheduler will use to receive bandwidth for its above-CIR offered-load. Values 1 to 8 (8 is the highest priority) Default 1 cir-weight <i>cir-weight</i> — 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 to 100 Default 0 cir-level <i>cir-level</i> — 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 to 8 (8 is the highest priority) Default 0

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] 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 than 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 the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.
Parameters	pir <i>pir</i> — The pir parameter configures the PIR rate of the scheduler in kb/s or it can be set to the maximum using the max keyword. Values 1 to 3200000000, max Default max cir <i>cir</i> — The cir parameter configures the CIR rate of the scheduler in kb/s or it can be set to the maximum using the max keyword. The sum keyword can also be used, which sets the CIR to the sum of child CIR Values. Values 1 to 3200000000, max, sum Default sum

8.4.2.3 Port Scheduler Policy Commands

port-scheduler-policy

Syntax	port-scheduler-policy <i>port-scheduler-name</i> [create] no port-scheduler-policy <i>port-scheduler-name</i>
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 (1 to 8)
- The cir associated with each priority level (1 to 8)

The orphan priority level (level 0) has no configuration parameters and cannot be overridden.

The **no** form of this 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.

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 effect when the port on which this **port-scheduler-policy** is configured is part of a LAG in **link** or **port-fair** mode. Similarly, when rates are configured as a **percent-rate**, this parameter is ignored.

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. 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.

CLI will enforce that mapping of levels to a group are contiguous. A user would not be able to add a priority level to a group unless the resulting set of priority levels is contiguous.

The **no** form of this 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.
 - create** — This keyword is mandatory when creating the specified group.

monitor-threshold

- Syntax** **monitor-threshold** *percent*
no monitor-threshold
- Context** config>qos>port-scheduler-policy
config>qos>port-scheduler-policy>group
- Description** This command defines the congestion monitoring threshold for the desired monitoring entity under the port-scheduler for per aggregate port-scheduler rate, per individual level, and per group that is aggregating multiple levels.
- The congestion threshold is specified in percentages of the configured PIR rate for the entity for which congestion monitoring is desired. For example, if the configured PIR rate for level 1 is 100,000 kb/s, and the monitoring threshold is set to 90%, then an event where the offered rate is >90,000 kb/s will be recorded. This event is shown as part of the cumulative count of congestion threshold exceeds since the last clearing of the counters.
- The **no** form of this command removes the congestion monitoring threshold.
- Default** no monitor-threshold
- Parameters** *percent* — Specifies the percent of the configured rate. If the offered rate exceeds the configured threshold, a counter monitoring the threshold will be increased.
- Values** 0 to 100

percent-rate

- Syntax** **percent-rate** *pir-percent* [**cir** *cir-percent*]
no percent-rate
- Context** config>qos>port-scheduler-policy>group

Description	The percent-rate command within the port scheduler policy group enables support for a policer's PIR and CIR rate to be configured as a percentage of the immediate parent root policer/arbitrator rate or the FP capacity. If the parent arbitrator 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 any time. 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/arbitrator 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/arbitrator rate or the FP capacity. Values Percentage ranging from 0.00 to 100.00. Default 100.00

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] 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-rate — Specifies PIR rates, in kilobits per second. Values 1 to 3200000000 cir cir-rate — Specifies CIR rates, in kilobits per second. Values 0 to 3200000000

level

Syntax	level <i>priority-level</i> rate <i>pir-rate</i> [cir <i>cir-rate</i>] group name [weight <i>weight</i>] [monitor-threshold

	<i>percent</i>] level <i>priority-level</i> percent-rate <i>pir-percent</i> [percent-cir <i>cir-percent</i>] group name [weight <i>weight</i>] [monitor-threshold <i>percent</i>] level <i>priority-level</i> rate <i>pir-rate</i> [cir <i>cir-rate</i>] [monitor-threshold <i>percent</i>] level <i>priority-level</i> percent-rate <i>pir-percent</i> [percent-cir <i>cir-percent</i>] [monitor-threshold <i>percent</i>] no level <i>priority-level</i>
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 zero 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 to 8 (8 is the highest priority) pir <i>pir</i> — Specifies the total bandwidth limits allocated to priority-level. Values 1 to 3200000000 max (kilobits per second (1,000 bits per second)) percent-rate <i>pir-percent</i> — Specifies the percent bandwidth limits allocated to priority-level. Values 0.01 to 100.00 max (kilobits per second) cir <i>cir</i> — The cir 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 0 to 3200000000 max (kilobits per second) The value given for kilobits per second is expressed in kilobits per second on a base 10 scale as 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). percent-cir <i>cir-percent</i> — Specifies the percent bandwidth limits allocated to priority-level. Values 0.01 to 100.00 max (kilobits per second) group <i>name</i> — specifies the existing group that the weighted scheduler group this level maps to. weight <i>weight</i> — Specifies the weight of the level within this weighted scheduler group. Values 1 to 100 Default 1

monitor-threshold *percent* — Specifies the percent of the configured rate. If the offered rate exceeds the configured threshold, a counter monitoring the threshold will be increased.

Values 0 to 100

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, when a scheduler policy is associated with a port or channel, the instance of the scheduler on the port automatically limits 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. This command can be used on channels only on the 7450 ESS and 7750 SR.

The max-rate command may be executed at any time 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. When 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 in kb/s.

Values 1 to 3200000000, max

percent *percent-rate* — Specifies the percent rate.

Values 0.01 to 100.00

orphan-override

Syntax **orphan-override** [*level priority-level*] [*weight weight*] [*cir-level cir-level*] [*cir-weight cir-weight*]
no orphan-override

Context	config>qos>port-scheduler-policy
Description	This command overrides 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 scheduler's 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 scheduler's within-CIR offered load. The within-CIR offered load is the amount of bandwidth the queue or scheduler 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 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 scheduler's 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 this 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 their above-CIR offered-load.

Values 1 to 8 (8 is the highest priority)

Default 1

weight *weight* — Defines the weight the orphan queues and schedulers will use in the above-CIR port priority level (defined by the level parameter).

Values 1 to 100

Default 1

cir-level *cir-level* — Defines the port priority the orphan queues and schedulers will use to receive bandwidth for their 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 to 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 to 100 (100 is the highest weight)

8.4.2.4 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 <i>src-name dst-name</i> — 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 <i>src-name dst-name</i> — 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.

8.4.2.5 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

scheduler-hierarchy

Syntax	scheduler-hierarchy
Context	show>qos
Description	This command displays scheduler-hierarchy information.

customer

Syntax	customer <i>customer-id site customer-site-name</i> [scheduler <i>scheduler-name</i>] [ingress egress] [detail]
Context	show>qos>scheduler-hierarchy
Description	This command displays the scheduler hierarchy per customer multiservice site.

Parameters

customer *customer-id* — Specifies the ID number associated with a particular customer.

Values 1 to 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 The following output is an example of customer scheduler hierarchy information, and [Table 80](#) describes the customer scheduler hierarchy fields.

Sample Output

```
*A:PE# show qos scheduler-hierarchy customer 1 site "s1"

=====
Scheduler Hierarchy - Customer 1 MSS s1
=====
Root (Ing)
| slot(1)
|--(S) : s1
|   |
|   |--(S) : Tier0Ingress:1->1/1/1->1
|   |   |
|   |   |--(Q) : 1->1/1/1->1 1-1
|   |   |
|   |   |--(Q) : 1->1/1/1->1 2-1
|   |   |
|   |   |--(Q) : 1->1/1/1->1 5-1
|   |
|
Root (Egr)
| slot(1)
|--(S) : s1
|   |
|   |--(Q) : 1->1/1/1->1
|   |
|

=====
*A:PE#
```

Table 80 Show QoS Scheduler-Hierarchy Customer 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.
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 multiservice site name.
Orphan Queues	Specifies the number of queues in an orphaned state.
Hierarchy	Displays the scheduler policy tree structure.

port

Syntax `port port-id queue-group queue-group-name [scheduler scheduler-name] [ingress | egress] [access | network] [instance instance-id] [detail]`

port *port-id* **vport** *name* **summary**
port *port-id* [**detail**]
port *port-id* **vport** *name* [**scheduler** *scheduler-name*] [**detail**]

Context show>qos>scheduler-hierarchy

Description This command displays scheduler hierarchy information per port.

Parameters *port-id* — Specifies the port ID.

<i>port-id</i>	<i>slot/mda/port</i> [<i>.channel</i>]		
	eth-sat-id	esat-id/slot/port	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-id.sub-port	
		pxc	keyword
		id	1 to 64
		sub-port	a, b

detail — Displays detailed information.

queue-group-name — Displays information about the specified queue group on the port.

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.

access — Displays scheduler-hierarchy statistics applied on an access port.

network — Displays scheduler-hierarchy statistics applied on a network port.

instance-id — Specifies the identification of a specific instance of the queue group.

vport *name* — Displays statistics for the named VPORT.

Output The following output is an example of port scheduler hierarchy information, and [Table 81](#) describes port scheduler hierarchy fields.

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#

```

Table 81 **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.

sap

Syntax **sap** *sap-id* [**scheduler** *scheduler-name*] [**ingress** | **egress**] [**detail**]
sap *sap-id* **encap-group** *group-name* [**scheduler** *scheduler-name*] [**member** *encap-id*]
 [**detail**]

Context show>qos>scheduler-hierarchy

Description This command displays the scheduler hierarchy per SAP.

Parameters **sap** *sap-id* — Specifies the SAP assigned to the service.

Values The following values apply to the 7750 SR and 7950 XRS:

<i>sap-id</i>	null	[<i>port-id</i> <i>lag-id</i>]
	dot1q	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1</i>
	qinq	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>
	<i>port-id</i>	<i>slot/mda/port</i> [. <i>channel</i>] <i>esat-id/slot/port</i>
	esat	keyword
	<i>id</i>	1 to 20
	<i>pxc-id.sub-port</i>	
	pxc	keyword
	<i>id</i>	1 to 64
	<i>sub-port</i>	a, b
	<i>lag-id</i>	<i>lag-id</i>
		<i>lag</i> keyword
		<i>id</i> 1 to 800
	<i>qtag1</i>	0 to 4094
	<i>qtag2</i>	*, 0 to 4094

Values The following values apply to the 7750 SR only:

<i>sap-id</i>	null	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i> <i>aps-id</i>]
	dot1q	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i> <i>aps-id</i>]: <i>qtag1</i>
	qinq	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>
	atm	[<i>port-id</i> <i>aps-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
	frame	[<i>port-id</i> <i>aps-id</i>]: <i>dlci</i>
	cisco-hdlc	<i>slot/mda/port.channel</i>
	cem	<i>slot/mda/port.channel</i>
	ima-grp	[<i>bundle-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
	port-id	<i>slot/mda/port</i> [<i>.channel</i>]
	bundle-id	<i>bundle-type-slot/mda.bundle-num</i>
	bundle	keyword
	type	ima, fr, ppp
	bundle-num	1 to 336
	bpgrp-id	<i>bpgrp-type-bpgrp-num</i>
	bpgrp	keyword
	type	ima, ppp
	bpgrp-num	1 to 2000
	aps-id	<i>aps-group-id</i> [<i>.channel</i>]
	aps	keyword
	group-id	1 to 64
	ccag-id	<i>ccag-id.path-id</i> [<i>cc-type</i>]: <i>cc-id</i>
	ccag	keyword
	id	1 to 8
	path-id	a, b
	cc-type	.sap-net, .net-sap
	cc-id	0 to 4094
	lag-id	<i>lag-id</i>
	lag	keyword
	id	1 to 800
	qtag1	0 to 4094
	qtag2	*, 0 to 4094
	vpi	NNI: 0 to 4095 UNI: 0 to 255
	vci	1, 2, 5 to 65535
	dlci	16 to 1022
	ipsec-id	<i>ipsec-id</i> [<i>private</i> <i>public</i>]: <i>tag</i>
	ipsec	keyword
	id	1 to 4
	tag	0 to 4094

Values The following values apply to the 7450 ESS only:

<i>sap-id</i>	null	[<i>port-id</i> <i>lag-id</i>]
	dot1q	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1</i>
	qinq	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>
	frame	[<i>port-id</i>]: <i>dlci</i>
	port-id	slot/mda/port[.channel]
	aps-id	aps-group-id[.channel]
	aps	keyword
	group-id	1 to 64
	ccag-id	ccag-id.path-id[cc-type]:cc-id
	ccag	keyword
	id	1 to 8
	path-id	a, b
	cc-type	.sap-net, .net-sap
	cc-id	0 to 4094
	lag-id	lag-id
	lag	keyword
	id	1 to 800
	qtag1	0. to 4094
	qtag2	*, 0 to 4094
	dlci	16 to 1022

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.

group-name — Specifies the name of the encap-group and can be up to 32 ASCII characters.

encap-id — Specifies the value of the single encap-id.

Values 1 to 16777215

detail — Displays detailed information.

Output The following output is an example of SAP scheduler hierarchy information, and [Table 82](#) describes the SAP scheduler hierarchy fields.

Sample Output

*A:PE# show qos scheduler-hierarchy sap 1/1/1:1 ingress detail

=====


```

Scheduler Hierarchy - Sap 1/1/1:1
=====
Ingress Scheduler Policy: sp1
-----
Legend :
(*) real-time dynamic value
(w) Wire rates
B   Bytes
-----

Root (Ing)
| slot(1)
|--(S) : s1
|   |   AdminPIR:200000          AdminCIR:30000 (sum)
|   |   Limit Unused Bandwidth: disabled
|   |
|   |   [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:200000
|   |
|   |   [As Parent]
|   |   Rate:200000
|   |   ConsumedByChildren:0
|   |
|--(Q) : 1->1/1/1:1->2 1-1
|   |   AdminPIR:100000      AdminCIR:20000
|   |   Parent Limit Unused Bandwidth: not-found
|   |   CBS:27648 B
|   |   MBS:125952 B
|   |   Depth:0 B
|   |   HighPlus Drop Tail:125952 B
|   |   High Drop Tail:125952 B
|   |   Low Drop Tail:110592 B
|   |   Exceed Drop Tail:98304 B
|   |
|   |   [Within CIR Level 2 Weight 1]
|   |   Assigned:20000       Offered:0
|   |   Consumed:0
|   |
|   |   [Above CIR Level 2 Weight 1]
|   |   Assigned:100000      Offered:0
|   |   Consumed:0
|   |
|   |   TotalConsumed:0
|   |   OperPIR:100000      OperCIR:20000
|   |   OperFIR:0
|   |
|   |   PktByteOffset:add 0*
|   |   OnTheWireRates:false

```

```

| | | ATMonTheWireRates:false
| | | LastMileOnTheWireRates:false
|
| | | --(Q) : 1->1/1/1:1->1 1-1
| | | AdminPIR:10000 AdminCIR:10000
| | | Parent Limit Unused Bandwidth: not-found
| | | CBS:15360 B
| | | MBS:66048 B
| | | Depth:0 B
| | | HighPlus Drop Tail:66048 B
| | | High Drop Tail:66048 B
| | | Low Drop Tail:58368 B
| | | Exceed Drop Tail:52224 B
|
| | | [Within CIR Level 6 Weight 1]
| | | Assigned:10000 Offered:0
| | | Consumed:0
|
| | | [Above CIR Level 6 Weight 1]
| | | Assigned:10000 Offered:0
| | | Consumed:0
|
| | | TotalConsumed:0
| | | OperPIR:10000 OperCIR:10000
| | | OperFIR:0
|
| | | PktByteOffset:add 0*
| | | OnTheWireRates:false
| | | ATMonTheWireRates:false
| | | LastMileOnTheWireRates:false

```

```

=====
*A:PE#
*A:PE# show qos scheduler-hierarchy sap 1/1/1:1 egress detail

```

```

=====
Scheduler Hierarchy - Sap 1/1/1:1
=====

```

```

Egress Scheduler Policy : spl
-----

```

```

Legend :
(*) real-time dynamic value
(w) Wire rates
B Bytes

```

```

-----
Root (Egr)
| slot(1)
| --(S) : s1
| | AdminPIR:200000 AdminCIR:10000 (sum)
| | Limit Unused Bandwidth: disabled
|
| | [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:200000
|
|
| | [As Parent]
| | Rate:200000
| | ConsumedByChildren:0
|
|
| | --(Q) : 1->1/1/1:1->1
| | AdminPIR:10000      AdminCIR:10000
| | Parent Limit Unused Bandwidth: not-found
| | CBS:15360 B
| | MBS:66048 B
| | Depth:0 B
| | HighPlus Drop Tail:66048 B
| | High Drop Tail:66048 B
| | Low Drop Tail:58368 B
| | Exceed Drop Tail:52224 B
|
| | [Within CIR Level 6 Weight 1]
| | Assigned:10000      Offered:0
| | Consumed:0
|
| | [Above CIR Level 6 Weight 1]
| | Assigned:10000      Offered:0
| | Consumed:0
|
| | TotalConsumed:0
| | OperPIR:10000      OperCIR:10000
|
| | PktByteOffset:add 0*
| | OnTheWireRates:false
| | ATMonTheWireRates:false
| | LastMileOnTheWireRates:false

```

=====

*A:PE#

Table 82 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.

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.
- Output** The following output is an example of scheduler name information.

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-policy

- Syntax** `scheduler-policy [scheduler-policy-name] [association | sap-ingress policy-id | sap-egress policy-id]`

	scheduler-policy <i>scheduler-policy-name</i> { ingress egress } queue-group <i>queue-grp-name</i>
Context	show>qos>scheduler-hierarchy
Description	Use this command to display scheduler policy information.
Parameters	<i>scheduler-policy-name</i> — The name of a scheduler policy configured in the config>qos>scheduler-policy context. association — Displays the associations related to the specified scheduler name. sap-ingress <i>policy-id</i> — Specifies the SAP ingress QoS policy information. sap-egress <i>policy-id</i> — Specifies the SAP egress QoS policy information. {ingress egress} — Specifies the direction to display and are mutually exclusive. <i>queue-grp-name</i> — Displays information for the specified queue group.
Output	The following output is an example of scheduler policy information, and Table 83 describes the customer scheduler hierarchy fields.

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#

Table 83 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.
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.

Table 83 Show QoS Scheduler-Policy Output Fields (Continued)

Label	Description
Multi Service Site	Specifies the multiservice site name.
Orphan Queues	Specifies the number of queues in an orphaned state.
Hierarchy	Displays the scheduler policy tree structure.

subscriber

Syntax **subscriber** *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress** | **egress**] [**detail**]
subscriber *sub-ident-string* **sla-profile** *sla-profile-name* **sap** *sap-id* [**scheduler** *scheduler-name*] [**detail**] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

Context show>qos>scheduler-hierarchy

Description This command displays the scheduler hierarchy per subscriber.

Parameters *sub-ident-string* — Displays the subscriber identification policy name.

scheduler-name — Displays the scheduler name.

ingress — Displays ingress SAP subscriber scheduler stats.

egress — Displays egress SAP subscriber scheduler stats.

detail — Displays detailed information.

sla-profile-name — Displays information for the specified sla-profile-name.

spi-sharing-type — Displays subscriber scheduler hierarchy information for sessions with the specified SLA Profile Instance sharing type.

Values per-group, per-session-ipoe, per-session-ppp

spi-sharing-id — Displays subscriber scheduler hierarchy information for sessions with the specified SLA Profile Instance sharing type and sharing identifier. The spi-sharing-id is:

- a group id when sharing type is per-group
- an internal PPP or IPoE sharing id when the grouping type is per-PPP or per-IPoE-session

Output The following output is an example of QoS scheduler hierarchy subscriber information, and [Table 84](#) describes the QoS scheduler hierarchy subscriber fields.

Sample Output

```
*A:PE# show qos scheduler-hierarchy subscriber "sub1"
```

```
=====
Scheduler Hierarchy - Subscriber sub1
```

```
=====
Ingress Scheduler Policy: basic-upstream-policy
Egress Scheduler Policy : basic-downstream-policy
-----

Root (Ing)
| slot(1)
|--(S) : basic-upstream
|   |
|   |--(Q) : Sub=sub1:basic-sla 200->1/1/1:1->2 1-1
|   |
|   |--(Q) : Sub=sub1:basic-sla 200->1/1/1:1->1 1-1
|   |
|   |
|   |

Root (Egr)
| slot(1)
|--(S) : basic-downstream (Port 1/1/1)
|   |
|   |--(Q) : Sub=sub1:basic-sla 200->1/1/1:1->1
|   |
|   |
|   |

=====
*A:PE#
*A:PE# show qos scheduler-hierarchy subscriber "sub1" detail

=====
Scheduler Hierarchy - Subscriber sub1
=====
Ingress Scheduler Policy: basic-upstream-policy
Egress Scheduler Policy : basic-downstream-policy
-----

Legend :
(*) real-time dynamic value
(w) Wire rates
B Bytes
-----

Root (Ing)
| slot(1)
|--(S) : basic-upstream
|   |
|   | AdminPIR:max AdminCIR:0(sum)
|   | Limit Unused Bandwidth: disabled
|   |
|   |
|   | [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:max
|   |
|   | [As Parent]
```

```

Rate:max
ConsumedByChildren:0

--(Q) : Sub=sub1:basic-sla 200->1/1/1:1->2 1-1
AdminPIR:1000000 AdminCIR:0
Parent Limit Unused Bandwidth: not-found
CBS:31457280 B
MBS:51904512 B
Depth:0 B
HighPlus Drop Tail:51904512 B
High Drop Tail:51904512 B
Low Drop Tail:46399488 B
Exceed Drop Tail:40894464 B

[Within CIR Level 0 Weight 1]
Assigned:0 Offered:0
Consumed:0

[Above CIR Level 1 Weight 1]
Assigned:1000000 Offered:0
Consumed:0

TotalConsumed:0
OperPIR:max OperCIR:0
OperFIR:0

PktByteOffset:add 0*
OnTheWireRates:false
ATMOnTheWireRates:false
LastMileOnTheWireRates:false

--(Q) : Sub=sub1:basic-sla 200->1/1/1:1->1 1-1
AdminPIR:10000 AdminCIR:0
Parent Limit Unused Bandwidth: not-found
CBS:10420224 B
MBS:20840448 B
Depth:0 B
HighPlus Drop Tail:20840448 B
High Drop Tail:20840448 B
Low Drop Tail:20840448 B
Exceed Drop Tail:16515072 B

[Within CIR Level 0 Weight 1]
Assigned:0 Offered:0
Consumed:0

[Above CIR Level 1 Weight 1]
Assigned:10000 Offered:0
Consumed:0

TotalConsumed:0
OperPIR:10000 OperCIR:0
OperFIR:0

PktByteOffset:add 0*
OnTheWireRates:false
ATMOnTheWireRates:false
LastMileOnTheWireRates:false

```

```

Root (Egr)
| slot(1)
|--(S) : basic-downstream (Port 1/1/1)
|       AdminPIR:max          AdminCIR:0 (sum)
|       Parent Limit Unused Bandwidth: not-found
|
|       AvgFrmOv:100.00
|       AdminPIR:max(w)      AdminCIR:0 (w)
|
|       [Within CIR Level 0 Weight 0]
|       Assigned:0 (w)      Offered:0 (w)
|       Consumed:0 (w)
|
|       [Above CIR Level 1 Weight 1]
|       Assigned:1000000(w) Offered:0 (w)
|       Consumed:0 (w)
|
|       TotalConsumed:0 (w)
|       OperPIR:1000000
|
|       [As Parent]
|       Rate:1000000
|       ConsumedByChildren:0
|
|--(Q) : Sub=sub1:basic-sla 200->1/1/1:1->1
|       AdminPIR:1000000    AdminCIR:0
|       Parent Limit Unused Bandwidth: not-found
|       AvgFrmOv:100.00
|       CBS:0 B
|       MBS:1253376 B
|       Depth:0 B
|       HighPlus Drop Tail:1253376 B
|       High Drop Tail:1253376 B
|       Low Drop Tail:1253376 B
|       Exceed Drop Tail:983040 B
|
|       [Within CIR Level 0 Weight 1]
|       Assigned:0          Offered:0
|       Consumed:0
|
|       [Above CIR Level 1 Weight 1]
|       Assigned:1000000    Offered:0
|       Consumed:0
|
|       TotalConsumed:0
|       OperPIR:max        OperCIR:0
|
|       PktByteOffset:add 0*
|       OnTheWireRates:false
|       ATMonTheWireRates:false
|       LastMileOnTheWireRates:false

```

=====

*A: PE#

Table 84 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.

scheduler-stats

Syntax **scheduler-stats**

Context show>qos

Description This command displays scheduler statistics information.

customer

Syntax **customer** *customer-id* **site** *customer-site-name* [**scheduler** *scheduler-name*] [**ingress** | **egress**]

Context show>qos>scheduler-stats

Description This command displays scheduler statistics customer information.

Parameters **customer** *customer-id* — Specifies the ID number associated with a particular customer.

Values 1 to 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 The following output is an example of SAP scheduler-stats customer information, and [Table 85](#) describes the SAP scheduler-stats customer fields.

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#
```

Table 85 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.

sap

Syntax **sap** *sap-id* [**scheduler** *scheduler-name*] [**ingress** | **egress**]
sap *sap-id* **encap-group** *group-name* [**scheduler** *scheduler-name*] [**member** *encap-id*]

Context show>qos>scheduler-stats

Description Display the scheduler stats per SAP.

Parameters **sap** *sap-id* — The port number and encapsulation value used to identify the SAP.

Values The following values apply to the 7750 SR and 7950 XRS:

<i>sap-id</i>	null	[<i>port-id</i> <i>lag-id</i>]
	dot1q	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1</i>
	qinq	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>
	lag-id	lag-id
		lag keyword
		id 1 to 800
	qtag1	0 to 4094
	qtag2	*, 0 to 4094
	vpi	NNI: 0 to 4095
		UNI: 0 to 255
	vci	1, 2, 5 to 65535
	dlci	16 to 1022
	ipsec-id	ipsec- <i>id</i> . [<i>private</i> <i>public</i>]: <i>tag</i>
		ipsec keyword
		id 1 to 4
		tag 0 to 4094

Values The following values apply to the 7750 SR only:

<i>sap-id</i>	null	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i> <i>aps-id</i>]
	dot1q	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i> <i>aps-id</i>]: <i>qtag1</i>
	qinq	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>
	atm	[<i>port-id</i> <i>aps-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
	frame	[<i>port-id</i> <i>aps-id</i>]: <i>dlci</i>
	cisco-hdlc	<i>slot/mda/port.channel</i>
	cem	<i>slot/mda/port.channel</i>
	ima-grp	[<i>bundle-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
	port-id	<i>slot/mda/port</i> [<i>.channel</i>]
	bundle-id	<i>bundle-type-slot/mda.bundle-num</i>
		bundle keyword
		type ima, fr, ppp
		bundle-num 1 to 336
	bpgrp-id	<i>bpgrp-type-bpgrp-num</i>
		bpgrp keyword
		type ima, ppp
		bpgrp-num 1 to 2000
	aps-id	<i>aps-group-id</i> [<i>.channel</i>]
		aps keyword
		group-id 1 to 64
	ccag-id	<i>ccag-id.path-id</i> [<i>cc-type</i>]: <i>cc-id</i>
		ccag keyword
		id 1 to 8
		path-id a, b
		cc-type .sap-net, .net-sap
		cc-id 0 to 4094
	eth-tunnel	<i>eth-tunnel-id</i> [<i>eth-tun-sap-id</i>]
		id: 1 to 1024
		eth-tun-sap- 0 to 4094
		id
	lag-id	lag-id
		lag keyword
		id 1 to 800
	qtag1	0 to 4094
	qtag2	*, 0 to 4094
	vpi	NNI: 0 to 4095
		UNI: 0 to 255
	vci	1, 2, 5 to 65535
	dlci	16 to 1022

ipsec-id ipsec-*id*.*[private | public]:tag*

Values The following values apply to the 7450 ESS:

<i>sap-id</i>	null	<i>[port-id lag-id]</i>
	dot1q	<i>[port-id lag-id]:qtag1</i>
	qinq	<i>[port-id lag-id]:qtag1.qtag2</i>
	frame	<i>[port-id]:dlci</i>
	port-id	slot/mda/port <i>[.channel]</i>
	aps-id	aps-group-id <i>[.channel]</i>
	aps	keyword
	group-id	1 to 64
	ccag-id	ccag-id. <i>path-id</i> <i>[cc-type]:cc-id</i>
	ccag	keyword
	<i>id</i>	1 to 8
	<i>path-id</i>	a, b]
	<i>cc-type</i>	.sap-net, .net-sap]
	<i>cc-id</i>	0 to 4094
	eth-tunnel	eth-tunnel-id <i>[:eth-tun-sap-id]</i>
	id:	1 to 1024
	eth-tun-sap-id	0 to 4094
	lag-id	lag- <i>id</i>
	lag	keyword
	<i>id</i>	1 to 800
	qtag1	0. to 4094
	qtag2	*, 0 to 4094
	dlci	16 to 1022

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.

group-name — Specifies the name of the encap-group and can be up to 32 ASCII characters.

encap-id — Specifies the value of the single encap-id.

Values 1 to 16777215

Output The following output is an example of scheduler-stats SAP information, and [Table 86](#) describes the scheduler-stats SAP fields.

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#
```

Table 86 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).

subscriber

Syntax **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**scheduler** *scheduler-name*] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
subscriber *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress** | **egress**]

Context show>qos>scheduler-stats

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.

sla-profile-name — Displays information for the specified sla-profile-name.

spi-sharing-type — Displays subscriber scheduler hierarchy information for sessions with the specified SLA Profile Instance sharing type.

Values per-group, per-session-ipoe, per-session-ppp

spi-sharing-id — Displays subscriber scheduler hierarchy information for sessions with the specified SLA Profile Instance sharing type and sharing identifier. The spi-sharing-id is:

- a group id when sharing type is per-group
- an internal PPP or IPoE sharing id when the sharing type is per-PPP or per-IPoE-session

Output The following output is an example of QoS scheduler-stats subscriber information, and [Table 87](#) describes the QoS scheduler-stats subscriber fields.

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#

```

Table 87 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.

agg-rate

Syntax **agg-rate**

Context show>qos

Description This command displays aggregate rate information.

customer

Syntax **customer** *customer-id* **site** *customer-site-name* [**egress**] [**detail**]

Context show>qos>agg-rate

Description This command displays the HQoS aggregate rate limit per customer multiservice site.

Parameters **customer** *customer-id* — Specifies the ID number associated with a particular customer.

Values 1 to 2147483647

site *customer-site-name* — The unique customer site name.

egress — Displays egress SAP customer scheduler stats.

detail — Displays detailed information.

Output The following output is an example of QoS customer aggregation rate output, and [Table 88](#) describes the QoS customer aggregation rate fields.

Sample Output

```
*A:PE1# show qos agg-rate customer 1 site "site1" egress
=====
Aggregate Rate Information - Customer 1 MSS site1
=====
Root (Egr)
| slot(1)
|   AdminRate           : 10000
|   OperRate            : 0
|   Limit Unused Bandwidth : disabled
|   OnTheWireRates       : false
|   LastMileOnTheWireRates : false
|
=====
*A:PE#
```

Table 88 Show QoS Aggregation Rate Customer Output Fields

Label	Description
AdminRate	
OperRate	
Limit Unused Bandwidth	Indicates whether the limit-unused-bandwidth command is enabled to protect against exceeding the aggregated bandwidth
OnTheWireRates	Indicates whether the displayed rates are on-the-wire rates.
LastMileOnTheWireRates	Indicates whether the displayed rates are on-the-wire rates for the last mile only.

port

Syntax **port** *port-id* **queue-group** *queue-group-name* [**egress**] [**access** | **network**] [**instance** *instance-id*] [**detail**]
port *port-id* **vport** *name* [**detail**]

Context show>qos>agg-rate

Description This command displays the HQoS aggregate rate limit per port or Vport.

Parameters *port-id* — Specifies the port ID.

<i>port-id</i>	<i>slot/mda/port [.channel]</i>		
	eth-sat-id	<i>esat-id/slot/port</i>	
		esat	keyword
		<i>id</i>	1 to 20
	pxc-id	<i>pxc-id.sub-port</i>	
		pxc	keyword
		<i>id</i>	1 to 64
		<i>sub-port</i>	a, b

queue-group *queue-group-name* — Displays information about the specified queue group on the port.

egress — Displays egress queue group information.

access — Displays HQoS aggregate rate limit information on an access port.

network — Displays HQoS aggregate rate limit information on a network port.

instance *instance-id* — Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

vport *name* — Displays HQoS aggregate rate limit information for the specified Vport.

detail — Displays detailed information.

Output The following output is an example of QoS port aggregation rate output, and [Table 89](#) describes the QoS port aggregation rate fields.

Sample Output

```
*A:PE# show qos agg-rate port 1/1/1 queue-group "qg1" instance 1 access egress
=====
Aggregate Rate Information - Queue Group Port 1/1/1
=====
-----
Egress Scheduler Policy :
Egress Queue-Group : qg1                               Instance-Id : 1
-----
Root (Egr)
| slot(1)
|   AdminRate           : 10000
|   OperRate            : 0
|   Limit Unused Bandwidth : disabled
|   OnTheWireRates       : false
|   LastMileOnTheWireRates : false
|
|-----
=====
```

*A: PE#

Table 89 Show QoS Aggregation Rate Port Output Fields

Label	Description
Egress Scheduler Policy	Indicates the name of the egress scheduler policy.
Egress Queue-Group	Indicates the queue group name.
Instance-Id	Indicates the identifier of a specific instance of the queue group.
AdminRate	
OperRate	
Limit Unused Bandwidth	Indicates whether the limit-unused-bandwidth command is enabled to protect against exceeding the aggregated bandwidth
OnTheWireRates	Indicates whether the displayed rates are on-the-wire rates.
LastMileOnTheWireRates	Indicates whether the displayed rates are on-the-wire rates for the last mile only.

sap

Syntax **sap** *sap-id* [**egress**] [**detail**]
sap *sap-id* **encap-group** *group-name* [**member** *encap-id*] [**detail**]

Context show>qos>agg-rate

Description This command displays the HQoS aggregate rate limit per SAP or encap group.

Parameters **sap** *sap-id* — The port number and encapsulation value used to identify the SAP.

Values

<i>sap-id</i>	null	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i> <i>aps-id</i>]
	dot1q	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i> <i>aps-id</i>]: <i>qtag1</i>
	qinq	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>
	atm	[<i>port-id</i> <i>aps-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i> <i>cp.conn-prof-id</i>]
	cp	keyword
	conn-prof-id	[1 to 8000]
	frame	[<i>port-id</i> <i>aps-id</i>]: <i>dlci</i>
	cisco-hdlc	<i>slot/mda/port.channel</i>
	cem	<i>slot/mda/port.channel</i>

ima-grp	<i>[bundle-id[:vpi/vci vpi vpi1.vpi2 cp.conn-prof-id]</i>
	cp keyword
	conn-prof-id [1to8000]
port-id	<i>slot/mda/port[.channel]</i>
bundle-id	<i>bundle-type-slot/mda.bundle-num</i>
	bundle keyword
	type ima, fr, ppp
	bundle-num 1 to 336
bpgrp-id	<i>bpgrp-type-bpgrp-num</i>
	bpgrp keyword
	type ima, ppp
	bpgrp-num 1 to 2000
aps-id	<i>aps-group-id[.channel]</i>
	aps keyword
	group-id 1 to 64
ccag-id	<i>ccag-id.path-id[cc-type]:cc-id</i>
	ccag keyword
	id 1 to 8
	path-id a, b
	cc-type .sap-net, .net-sap
	cc-id 0 to 4094
eth-tunnel	<i>eth-tunnel-id[:eth-tun-sap-id]</i>
	id: 1 to 128
	eth-tun-sap-id 0 to 4094
lag-id	lag-id
	lag keyword
	id 1 to 800
pw-id	<i>pw-id</i>
	pw keyword
	id 1to 10239
qtag1	0 to 4094
qtag2	*, null, 0 to 4094
vpi	NNI: 0 to 4095
	UNI: 0 to 255
vci	1, 2, 5 to 65535
dlci	16 to 1022
tunnel-id	<i>tunnel-id.[private public]:tag</i>
	tunnel keyword

id 1 to 16
tag 0 to 4094

egress — Displays egress SAP customer scheduler stats.

group-name — Specifies the name of the encap-group and can be up to 32 ASCII characters.

encap-id — Specifies the value of the single encap-id.

Values 1 to 16777215

detail — Displays detailed information.

Output The following output is an example of HQoS aggregate rate limit per SAP output, and [Table 90](#) describes the HQoS aggregate rate limit per SAP fields.

Sample Output

```
*A:PE# show qos agg-rate sap 1/1/1:1 egress
=====
Aggregate Rate Information - Sap 1/1/1:1
=====
-----
Root (Egr)
| slot(1)
|   AdminRate           : 10000
|   OperRate            : 0
|   Limit Unused Bandwidth : disabled
|   OnTheWireRates       : false
|   LastMileOnTheWireRates : false
|
=====
*A:PE#
```

Table 90 Show HQoS Aggregate Rate Limit per SAP Output Fields

Label	Description
AdminRate	
OperRate	
Limit Unused Bandwidth	Indicates whether the limit-unused-bandwidth command is enabled to protect against exceeding the aggregated bandwidth
OnTheWireRates	Indicates whether the displayed rates are on-the-wire rates.
LastMileOnTheWireRates	Indicates whether the displayed rates are on-the-wire rates for the last mile only.

subscriber

- Syntax** **subscriber** *sub-indent-string* [**egress**] [**detail**]
- Context** show>qos>agg-rate
- Description** This command displays the HQoS aggregate rate limit per subscriber.
- Parameters** *sub-indent-string* — Specifies the subscriber identification string of the subscriber.
egress — Displays egress SAP customer scheduler stats.
detail — Displays detailed information.
- Output** The following output is an example of HQoS aggregate rate limit per subscriber, and [Table 91](#) describes the HQoS aggregate rate limit per subscriber fields.

Sample Output

```
*A:PE1# show qos agg-rate subscriber "sub1" egress
=====
Aggregate Rate Information - Subscriber sub1
=====
-----
Root (Egr)
| slot(1)
|   AdminRate           : 1000
|   OperRate            : 0
|   Limit Unused Bandwidth : disabled
|   OnTheWireRates       : false
|   LastMileOnTheWireRates : false
|
|-- (Q) : Sub=sub1:basic-sla 1->1/1/1:1->1 (Port 1/1/1)
|
=====
*A:PE#
```

Table 91 Show HQoS Aggregate Rate Limit per Subscriber Fields

Label	Description
AdminRate	
OperRate	
Limit Unused Bandwidth	Indicates whether the limit-unused-bandwidth command is enabled to protect against exceeding the aggregated bandwidth
OnTheWireRates	Indicates whether the displayed rates are on-the-wire rates.
LastMileOnTheWireRates	Indicates whether the displayed rates are on-the-wire rates for the last mile only.

port-scheduler-policy

Syntax	port-scheduler-policy [<i>port-scheduler-policy-name</i>] [association] port-scheduler-policy <i>port-scheduler-policy-name</i> network-policy <i>network-queue-policy-name</i> port-scheduler-policy <i>port-scheduler-policy-name</i> sap-egress <i>policy-id</i> port-scheduler-policy <i>port-scheduler-policy-name</i> scheduler-policy <i>scheduler-policy-name</i> port-scheduler-policy <i>port-scheduler-policy-name</i> scheduler-policy <i>scheduler-policy-name</i> sap-egress <i>policy-id</i>
Context	show>qos
Description	This command displays port-scheduler policy information
Parameters	<i>port-scheduler-policy-name</i> — Displays information for the specified existing port scheduler policy. association — Displays associations related to the specified port scheduler policy. network-policy <i>network-queue-policy-name</i> — Displays information for the specified existing network queue policy. sap-egress <i>policy-id</i> — Displays information for the specified existing SAP egress policy. scheduler-policy <i>scheduler-policy-name</i> — Displays information for the specified existing scheduler policy.
Output	The following output is an example of QoS port scheduler policy information, and Table 92 describes the QoS port scheduler policy fields.

Sample Output

```
*A:Dut-R# show qos port-scheduler-policy p1
=====
QoS Port Scheduler Policy
=====
Policy-Name           : p1
Max Rate              : max
Lv11 PIR              : max
Lv12 PIR              : max
Lv13 PIR              : max
Lv14 PIR              : max
Lv15 PIR              : max
Lv16 PIR              : max
Lv17 PIR              : max
Lv18 PIR              : max
Orphan Lvl           : default
Orphan CIR-Lvl       : default
Last changed          : 05/21/2007 10:39:15
Lv11 CIR              : max
Lv12 CIR              : max
Lv13 CIR              : max
Lv14 CIR              : max
Lv15 CIR              : max
Lv16 CIR              : max
Lv17 CIR              : max
Lv18 CIR              : max
Orphan Weight         : default
Orphan CIR-Weight    : default
=====
QoS Port Scheduler Policy
=====
Policy-Name           : p1
=====
```

Associations

```
-----
- Port : 5/1/1
=====
```

```
*A:Dut-R#
```

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
Dist LAG Rate   : True
Group           : test
Group PIR       : max
Group PIR Percent : 20.00
Max Rate Percent : 30.00
Last changed    : 07/16/2014 21:31:51
Group CIR       : max
Group CIR Percent : 20.00

Lv11 PIR        : max
Lv11 PIR Percent : 10.00
Lv12 PIR        : max
Lv12 PIR Percent : 20.00
Lv13 PIR        : max
Lv13 PIR Percent : 30.00
Lv14 PIR        : max
Lv14 PIR Percent : 40.00
Lv15 PIR        : max
Lv15 PIR Percent : 50.00
Lv16 PIR        : max
Lv16 PIR Percent : 60.00
Lv17 PIR        : max
Lv17 PIR Percent : 70.00
Lv18 PIR        : max
Lv18 PIR Percent : 80.00
Orphan Lvl      : default
Orphan CIR-Lvl  : default

Lv11 CIR        : max
Lv11 CIR Percent : 10.00
Lv12 CIR        : max
Lv12 CIR Percent : 20.00
Lv13 CIR        : max
Lv13 CIR Percent : 30.00
Lv14 CIR        : max
Lv14 CIR Percent : 40.00
Lv15 CIR        : max
Lv15 CIR Percent : 50.00
Lv16 CIR        : max
Lv16 CIR Percent : 60.00
Lv17 CIR        : max
Lv17 CIR Percent : 70.00
Lv18 CIR        : max
Lv18 CIR Percent : 80.00
Orphan Weight   : default
Orphan CIR-Weight : default
=====
Part of show port Output
-----
Egr Port Sched Override
-----
Max Rate      : max*
Lv11 PIR      : max*
Lv11 PIR Percent : 10.00
Lv12 PIR      : max*
Lv12 PIR Percent : 20.00
Lv13 PIR      : max*
Lv13 PIR Percent : 30.00
Lv14 PIR      : max*
Lv14 PIR Percent : 40.00
Lv15 PIR      : max*
Lv15 PIR Percent : 50.00
Lv16 PIR      : max*
Lv16 PIR Percent : 60.00
Lv17 PIR      : max*
Lv17 PIR Percent : 70.00
Lv18 PIR      : max*
Lv18 PIR Percent : 80.00
* means the value is inherited

Max Rate Percent : 50.00
Lv11 CIR         : max*
Lv11 CIR Percent : 10.00
Lv12 CIR         : max*
Lv12 CIR Percent : 20.00
Lv13 CIR         : max*
Lv13 CIR Percent : 30.00
Lv14 CIR         : max*
Lv14 CIR Percent : 40.00
Lv15 CIR         : max*
Lv15 CIR Percent : 50.00
Lv16 CIR         : max*
Lv16 CIR Percent : 60.00
Lv17 CIR         : max*
Lv17 CIR Percent : 70.00
Lv18 CIR         : max*
Lv18 CIR Percent : 80.00

```

Table 92 Show QoS Port Scheduler Output Fields

Label	Description
Policy Name	Displays the port scheduler policy name.

Table 92 Show QoS Port Scheduler Output Fields (Continued)

Label	Description (Continued)
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.
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.

8.4.2.6 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 The following values apply to the 7750 SR and 7950 XRS:

<i>sap-id</i>	null	[<i>port-id</i> <i>lag-id</i>]
	dot1q	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1</i>
	qinq	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>
	atm	[<i>port-id</i> <i>aps-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
	frame	[<i>port-id</i> <i>aps-id</i>]: <i>dlci</i>
	cisco-hdlc	<i>slot/mda/port.channel</i>
	cem	<i>slot/mda/port.channel</i>
	ima-grp	[<i>bundle-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
	port-id	<i>slot/mda/port</i> [. <i>channel</i>]
	lag-id	lag-id
		lag keyword
		id 1 to 800
	qtag1	0 to 4094
	qtag2	*, 0 to 4094

Values The following values apply to the 7750 SR:

<i>sap-id</i>	null	<i>[port-id bundle-id bpgrp-id lag-id aps-id]</i>
	dot1q	<i>[port-id bundle-id bpgrp-id lag-id aps-id]:qtag1</i>
	qinq	<i>[port-id bundle-id bpgrp-id lag-id]:qtag1.qtag2</i>
	atm	<i>[port-id aps-id][:vpi/vci vpi vpi1.vpi2]</i>
	frame	<i>[port-id aps-id]:dlci</i>
	cisco-hdlc	<i>slot/mda/port.channel</i>
	cem	<i>slot/mda/port.channel</i>
	ima-grp	<i>[bundle-id[:vpi/vci vpi vpi1.vpi2]</i>
	port-id	<i>slot/mda/port[.channel]</i>
	bundle-id	<i>bundle-type-slot/mda.bundle-num</i>
		bundle keyword
		type ima, fr, ppp
		bundle-num 1 to 336
	bpgrp-id	<i>bpgrp-type-bpgrp-num</i>
		bpgrp keyword
		type ima, ppp
		bpgrp-num 1 to 2000
	aps-id	<i>aps-group-id[.channel]</i>
		aps keyword
		group-id 1 to 64
	ccag-id	<i>ccag-id.path-id[cc-type]:cc-id</i>
		ccag keyword
		id 1 to 8
		path-id a, b
		cc-type .sap-net, .net-sap
		cc-id 0 to 4094
	lag-id	<i>lag-id</i>
		lag keyword
		id 1 to 800
	qtag1	0 to 4094
	qtag2	*, 0 to 4094
	vpi	NNI: 0 to 4095 UNI: 0 to 255
	vci	1, 2, 5 to 65535
	dlci	16 to 1022
	ipsec-id	<i>ipsec-id.[private public]:tag</i>
		ipsec keyword
		id 1 to 4
		tag 0 to 4094

Values The following values apply to the 7450 ESS:

<i>sap-id:</i>	null	[<i>port-id</i> <i>lag-id</i>]
	dot1q	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1</i>
	qinq	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>
	frame	[<i>port-id</i>]: <i>dlci</i>
	port-id	slot/mda/port[.channel]
	aps-id	aps-group-id[.channel]
	aps	keyword
	group-id	1 to 64
	ccag-id	ccag-id.path-id[cc-type]:cc-id
	ccag	keyword
	id	1 to 8
	path-id	a, b
	cc-type	.sap-net, .net-sap
	cc-id	0 to 4094
	lag-id	lag-id
	lag	keyword
	id	1 to 800
	qtag1	0. to 4094
	qtag2	*, 0 to 4094
	dlci	16 to 1022

scheduler-name — The name of the scheduler.

Values Valid names consist of any string up to 32 characters, 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.

9 Slope QoS Policies

9.1 Overview

Default buffer pools exist (logically) at the port, FP or MDA, and node levels. Each physical port has three associated pool objects:

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

Each FP has five associated pool objects:

- Access egress pool
- Access ingress pool
- Acc-Ing MC Path Mgmt
- Network ingress pool
- Network egress pool

By default, each pool is associated with slope-policy default, which has the slopes within the pool shutdown.

Slope policies are also applied when WRED is applied per queue and to the associated megapools; see [WRED Per Queue](#).

9.2 Basic Configurations

A basic slope QoS policy must conform to the following:

- Each slope policy must have a unique policy ID.
- Highplus slope, high slope, low slope, and exceed slope must be shut down (default).
- Default values can be modified but parameters cannot be deleted.

9.2.1 Creating a Slope QoS Policy

Configuring and applying slope policies is optional. If no slope policy is explicitly applied to a pool, 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 highplus slope for the highplus Random Early Detection (RED) slope graph.
- The high slope for the high Random Early Detection (RED) slope graph.
- The low slope for the low Random Early Detection (RED) slope graph.
- The exceed slope for the exceed 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 [create]
description description-string
exceed-slope
    max-avg percent
    max-prob percent
    start-avg percent
    no shutdown
high-slope
    max-avg percent
    max-prob percent
    start-avg percent
```

```

        no shutdown
    highplus-slope
        max-avg percent
        max-prob percent
        start-avg percent
        no shutdown
    low-slope
        max-avg percent
        max-prob percent
        start-avg percent
        no shutdown
    time-average-factor value

```

The following displays the slope policy configuration:

```

A:ALA-7>config>qos# info
#-----
echo "QoS Slope/Queue Policies Configuration"
#-----
...
    slope-policy "slopePolicy1" create
        description "Test"
        highplus-slope
            no shutdown
        exit
        high-slope
            no shutdown
        exit
        low-slope
            no shutdown
        exit
        exceed-slope
            no shutdown
        exit
    exit
...
#-----
A:ALA-7>config>qos#

```

9.2.2 Applying Slope Policies

Apply slope policies to the following entities:

9.2.2.1 Ports

The following CLI syntax examples may be used to apply slope policies to port pools:

CLI Syntax: config>port>access>ingress>pool>slope-policy name
config>port>access>egress>pool>slope-policy name
config>port>network>egress>pool>slope-policy name

9.2.2.2 FP

The following CLI syntax examples may be used to apply slope policies to FP pools:

CLI Syntax: config>card>mda>access>ingress>pool>slope-policy name
config>card>mda>access>egress>pool>slope-policy name
config>card>mda>network>ingress>pool>slope-policy name
config>card>mda>network>egress>pool>slope-policy name

Although these commands are applied under the MDA CLI tree, the policy is applied to the FP. If different policies are applied to each MDA, the last policy applied takes precedence.

For the MC Path Mgmt, the slope policy is configured within a multicast management bandwidth policy that is then applied under the FP ingress multicast-path-management.

CLI Syntax: config>mcast-mgmt>bw-plcy> mcast-pool percent-of-total
<percent-of-buffers> resv-cbs <percent-of-pool> slope-policy <policy-name>

9.2.3 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. [Table 93](#) lists default policy parameters.

Table 93 Slope Policy Defaults

Parameter	Default
Description	"Default slope policy"
Highplus-slope	

Table 93 **Slope Policy Defaults (Continued)**

Parameter		Default
	shutdown	shutdown
	start-avg	85
	max-avg	100
	max-prob	80
High-slope		
	shutdown	shutdown
	start-avg	70
	max-avg	90
	max-prob	80
Low-slope		
	shutdown	shutdown
	start-avg	50
	max-avg	75
	max-prob	80
Exceed-slope		
	shutdown	shutdown
	start-avg	30
	max-avg	55
	max-prob	80
Time-average-factor		7

The following output displays the default configuration:

```
A:ALA-7>config>qos>slope-policy# info detail
-----
description "Default slope policy."
highplus-slope
  shutdown
  start-avg 85
  max-avg 100
  max-prob 80
exit
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
exceed-slope
    shutdown
    start-avg 30
    max-avg 55
    max-prob 80
exit
time-average-factor 7
-----
A:ALA-7>config>qos>slope-policy#
```

9.3 Service Management Tasks

This section discusses slope QoS policy service management tasks:

9.3.1 Deleting QoS Policies

A slope policy is associated by default with XMA and MDAs 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 of the places where it is applied.

```
A:ALA-7>config>qos# no slope-policy slopePolicy1
MINOR: QoS #1902 Slope policy has references
A:ALA-7>config>qos#
```

9.3.1.1 Ports

The following CLI syntax examples can be used to remove slope policies from port pools:

CLI Syntax:

```
config>port>access>ingress>pool# no slope-policy name
config>port>access>egress>pool# no slope-policy name
config>port>network>egress>pool# no slope-policy name
```

9.3.1.2 FP

The following CLI syntax examples can be used to remove slope policies from FP pools:

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>ingress>pool# no slope-policy
name
config>card>mda>network>egress>pool# no slope-policy
name
```

Although some commands are applied under the MDA CLI tree, the policy is applied to the FP. If different policies are applied to each MDA, the last policy applied takes precedence.

For the MC Path Mgmt, the mcast-pool can be removed from the bandwidth policy or the bandwidth-policy can be removed from under the FP ingress multicast-path-management.

CLI Syntax: `config>mcast-mgmt>bw-plcy# no mcast-pool`
 `config>card>fp>ingress>mcast-mgmt# no bandwidth-policy`

9.3.1.3 Removing 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`

9.3.2 Copying and Overwriting QoS Policies

An existing slope policy can be copied, renamed with a new policy ID value, or used to 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]`

9.3.3 Editing QoS Policies

Existing policies and entries in the CLI or NMS can be edited. 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, then write over the original policy.

9.4 Slope QoS Policy Command Reference

9.4.1 Command Hierarchies

9.4.1.1 Configuration Commands

```
config
  — qos
    — slope-policy name [create]
    — no slope-policy name
      — description description-string
      — no description
      — [no] exceed-slope
        — max-avg percent
        — no max-avg
        — max-prob percent
        — no max-prob
        — [no] shutdown
        — start-avg percent
        — no start-avg
      — [no] high-slope
        — max-avg percent
        — no max-avg
        — max-prob percent
        — no max-prob
        — [no] shutdown
        — start-avg percent
        — no start-avg
      — [no] highplus-slope
        — max-avg percent
        — no max-avg
        — max-prob percent
        — no max-prob
        — [no] shutdown
        — start-avg percent
        — no start-avg
      — [no] low-slope
        — max-avg percent
        — no max-avg
        — max-prob percent
        — no max-prob
        — [no] shutdown
        — start-avg percent
        — no start-avg
      — time-average-factor value
      — no time-average-factor
```

9.4.1.2 Operational Commands

```
config
  — qos
    — copy slope-policy src-name dst-name [overwrite]
```

9.4.1.3 Show Commands

```
show
  — qos
    — slope-policy [slope-policy-name] [detail]
```

9.4.2 Command Descriptions

9.4.2.1 Configuration Commands

9.4.2.1.1 Generic Commands

description

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

shutdown

Syntax	[no] shutdown
Context	config>qos>slope-policy>exceed-slope config>qos>slope-policy>highplus-slope 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

9.4.2.2 Slope Policy QoS Commands

slope-policy

Syntax	slope-policy <i>name</i> [create] no slope-policy <i>name</i>
Context	config>qos
Description	This command enters the context to configure a QoS slope policy.
Default	slope-policy "default"
Parameters	<i>name</i> — The name of the slope policy. Values Valid names consist of any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

exceed-slope

Syntax	[no] exceed-slope
Context	config>qos>slope-policy

Description The **exceed-slope** context contains the commands and parameters for defining the exceed Random Early Detection (RED) slope graph. Each egress buffer pool supports an exceed RED slope for managing access to the shared portion of the buffer pool for exceed-profile packets.

The **exceed-slope** parameters can be changed at any time and the affected buffer pool exceed RED slopes are adjusted appropriately.

The **no** form of this command restores the exceed slope configuration commands to the default values. If the leaf commands within **exceed-slope** are set to the default parameters, the **exceed-slope** node will not appear in save config and show config output unless the detail parameter is present.

max-avg

Syntax **max-avg** *percent*
no max-avg

Context config>qos>slope-policy>exceed-slope
config>qos>slope-policy>highplus-slope
config>qos>slope-policy>high-slope
config>qos>slope-policy>low-slope

Description Sets the exceed, low, high, or highplus 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 100 - Highplus slope default is 100% buffer utilization before discard probability is 1.
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.
max-avg 55 — Exceed slope default is 55% 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 one. 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 to 100

max-prob

Syntax	max-prob <i>percent</i> no max-prob
Context	config>qos>slope-policy>exceed-slope config>qos>slope-policy>highplus-slope config>qos>slope-policy>high-slope config>qos>slope-policy>low-slope
Description	Sets the exceed, low, high, or highplus 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
Parameters	<i>percent</i> — The maximum drop probability percentage corresponding to the max-avg, expressed as a decimal integer. Values 0 to 100

start-avg

Syntax	start-avg <i>percent</i> no start-avg
Context	config>qos>slope-policy>exceed-slope config>qos>slope-policy>highplus-slope config>qos>slope-policy>high-slope config>qos>slope-policy>low-slope
Description	This command sets the exceed, low, high, or highplus 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.
Default	start-avg 85 - Highplus slope default is 85% buffer utilization before discard probability starts to increase above zero. start-avg 70 — High slope default is 70% buffer utilization before discard probability starts to increase above zero.

start-avg 50 — Low slope default is 50% buffer utilization before discard probability starts to increase above zero.

start-avg 30 — Exceed slope default is 30% buffer utilization before discard probability starts to increase above zero.

Parameters *percent* — The percentage of the shared buffer space for the buffer pool at which point the drop probability starts to increase above zero. The value entered must be less than or equal to the current setting of **max-avg**. If the entered value is greater than the current value of **max-avg**, an error will occur and no change will take place.

Values 0 to 100

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 Random Early Detection (RED) slope graph. Each buffer pool supports a high RED slope for managing access to the shared portion of the buffer pool for in-profile packets.

The **high-slope** parameters can be changed at any time and the affected buffer pool high 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.

highplus-slope

Syntax [no] highplus-slope

Context config>qos>slope-policy

Description The **highplus-slope** context contains the commands and parameters for defining the highplus Random Early Detection (RED) slope graph. Each buffer pool supports a highplus RED slope for managing access to the shared portion of the buffer pool for inplus-profile packets.

The **highplus-slope** parameters can be changed at any time and the affected buffer pool highplus RED slopes will be adjusted appropriately.

The **no** form of this command restores the highplus slope configuration commands to the default values. If the commands within **highplus-slope** are set to the default parameters, the **highplus-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 Random Early Detection (RED) slope graph. Each buffer pool supports a low RED slope for managing access to the shared portion of the buffer pool for low out-of-profile packets. The low-slope parameters can be changed at any time and the affected buffer pool low 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.

time-average-factor

Syntax	time-average-factor <i>value</i> 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	time-average-factor 7
Parameters	<i>value</i> — 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 to 15

9.4.2.3 Operational Commands

copy

Syntax	copy slope-policy <i>src-name dst-name</i> [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. <pre>ALA-7>config>qos# copy slope-policy default sp1 MINOR: CLI Destination "sp1" exists - use {overwrite}. ALA-7>config>qos#overwrite</pre>

9.4.2.4 Show Commands

slope-policy

Syntax	slope-policy [<i>slope-policy-name</i>] [detail]
Context	show>qos
Description	This command displays slope policy information.
Parameters	<i>slope-policy-name</i> — The name of the slope policy. detail — Displays detailed information about the slope policy.
Output	The following output is an example of slope QoS policy information, and Table 94 describes slope QoS policy output fields.

Sample Output

```
*A:PE# show qos slope-policy "slope1"
=====
QoS Slope Policy
=====
Policy          : slope1
Description     : (Not Specified)
Time Avg       : 7
-----
HighPlus Slope Parameters
-----
Start Avg      : 75                      Admin State : Enabled
Max Avg       : 100                     Max Prob.  : 100
-----
High Slope Parameters
-----
Start Avg      : 50                      Admin State : Enabled
Max Avg       : 65                      Max Prob.  : 100
-----
Low Slope Parameters
-----
Start Avg      : 30                      Admin State : Enabled
Max Avg       : 40                      Max Prob.  : 100
-----
Exceed Slope Parameters
-----
Start Avg      : 10                      Admin State : Enabled
Max Avg       : 25                      Max Prob.  : 100
=====
*A:PE#
*A:PE# show qos slope-policy "slope1" detail
=====
QoS Slope Policy
=====
Policy          : slope1
Description     : (Not Specified)
Time Avg       : 7
-----
HighPlus Slope Parameters
-----
Start Avg      : 75                      Admin State : Enabled
Max Avg       : 100                     Max Prob.  : 100
-----
High Slope Parameters
-----
Start Avg      : 50                      Admin State : Enabled
Max Avg       : 65                      Max Prob.  : 100
-----
Low Slope Parameters
-----
Start Avg      : 30                      Admin State : Enabled
Max Avg       : 40                      Max Prob.  : 100
-----
Exceed Slope Parameters
-----
Start Avg      : 10                      Admin State : Enabled
Max Avg       : 25                      Max Prob.  : 100
-----
Associations
-----
```

```

Object Type Object Id      Application      Pool
-----
No Matching Entries
-----
Named-Pool Associations
-----
Policy-Name                Pool-Name
-----
No Matching Entries
-----
WRED Slot/FP Associations
-----
Slot      FP
-----
No Matching Entries
-----
WRED SAP Queue Associations
-----
Policy-Id  Queue-Id
-----
No Matching Entries
-----
WRED QoS Queue-Group Associations
-----
Queue-Group                Queue-Id
-----
No Matching Entries
=====
*A:PE#

```

Table 94 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 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 one, expressed as a decimal integer.
Admin State	Enabled The administrative status of the RED slope is enabled. Disabled The administrative status of the RED slope is disabled.

Table 94 **Show QoS Slope Policy Output Fields (Continued)**

Label	Description
Max Prob.	Specifies the RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one.

10 Shared-Queue QoS Policies

10.1 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. Shared-queuing can add latency. Network planners should consider these restrictions while trying to scale services on the router.

10.2 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.

10.2.1 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.

10.2.2 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 35](#) shows 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 to the switch fabric.

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 35 Unicast Service Queue Mapping to Multiple Destination-Based Hardware Queues




Fig_22

10.2.2.1 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 36](#) shows shared unicast queuing. SAP or SLA profile instance multipoint queuing is not affected by enabling shared queuing.

Figure 36 Unicast Service Queuing With Shared Queuing Enabled





Fig_23

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, as shown in [Figure 37](#).

Enabling shared queuing may affect ingress performance due to double packet processing through the service and shared queues.

Figure 37 Multipoint Queue Behavior with Shared Queuing Enabled



10.2.2.2 Ingress Multipoint Shared Queuing

Ingress multipoint shared queuing is a variation of the unicast shared queuing defined in [Ingress Shared Queuing](#). With ingress multipoint shared queuing, ingress unicast service queues are mapped one-for-one with hardware queues and unicast packets traverse the ingress forwarding plane twice. In addition, 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, 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 reduce the overall subscriber density on the ingress forwarding plane. Multipoint shared queuing eliminates the multipoint queues, sparing hardware queues for better subscriber density. [Figure 38](#) shows 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 limited 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 38 Multipoint Shared Queuing Using First Pass Unicast Queues



Multipoint shared queuing cannot be enabled on the following services:

- Epipe
- Apipe (applies to the 7750 SR and 7950 XRS)
- Fpipe (applies to the 7750 SR and 7950 XRS)
- Ipipe
- Routed CO

For information about the tasks and commands necessary to access the CLI and to configure and maintain the router, refer to CLI Usage chapter in the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Basic System Configuration Guide*.

10.3 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.

10.3.1 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
    drop-tail low percent-reduction-from-mbs
    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
  drop-tail
  low
  percent-reduction-from-mbs 20
  exit
exit
exit
-----
A:ALA-48>config>qos>shared-queue#
```

10.3.1.1 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:

10.3.1.1.1 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#
```

10.3.1.1.2 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#
```

10.3.1.1.3 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
  exit
-----
A:SR>config>service#
```

10.3.1.1.4 VPRN Services

Use the following CLI syntax to apply QoS policies to ingress VPRN SAPs on the 7750 SR and 7950 XRS:

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-cel" create
    address 192.168.0.0/24
    sap 1/1/10:1 create
      ingress
        qos 1 multipoint-shared
      exit
    exit
  exit
  no shutdown
exit
-----
A:SR7>config>service#
```

10.3.2 Default Shared Queue Policy Values

The following shared queue policies are created by the system and cannot be deleted.

Table 95 **Default Shared Queue Policies**

Policy-ID	Description
default	Default Shared Queue Policy
policer-output-queues	Default Policer Output Shared Queue Policy
egress-pbr-ingress-queues	Egress PBR Ingress Shared Queue Policy

Creating additional shared queue policies is not permitted.

The following output displays the default configuration:

```
*A:ian1>config>qos>shared-queue# info detail
-----
description "Default Shared Queue Policy"
queue 1 auto-expedite create
    rate 100 cir 0 fir 0
    mbs 50
    cbs 1
    no pool
    drop-tail
    low
        percent-reduction-from-mbs default
    exit
exit
queue 2 auto-expedite create
    rate 100 cir 25 fir 0
    mbs 50
    cbs 3
    no pool
    drop-tail
    low
        percent-reduction-from-mbs default
    exit
exit
queue 3 auto-expedite create
    rate 100 cir 25 fir 0
    mbs 50
    cbs 10
    no pool
    drop-tail
    low
        percent-reduction-from-mbs default
    exit
exit
queue 4 auto-expedite create
    rate 100 cir 25 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
        percent-reduction-from-mbs default
    exit
exit
queue 5 auto-expedite create
    rate 100 cir 100 fir 0
    mbs 50
    cbs 10
    no pool
    drop-tail
    low
        percent-reduction-from-mbs default
    exit
exit
queue 6 auto-expedite create
```

```
        rate 100 cir 100 fir 0
        mbs 50
        cbs 10
        no pool
        drop-tail
        low
        percent-reduction-from-mbs default
        exit
    exit
exit
queue 7 auto-expedite create
    rate 100 cir 10 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
    exit
exit
queue 8 auto-expedite create
    rate 100 cir 10 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
    exit
exit
queue 9 auto-expedite multipoint create
    rate 100 cir 0 fir 0
    mbs 50
    cbs 1
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
    exit
exit
queue 10 auto-expedite multipoint create
    rate 100 cir 25 fir 0
    mbs 50
    cbs 3
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
    exit
exit
queue 11 auto-expedite multipoint create
    rate 100 cir 25 fir 0
    mbs 50
    cbs 10
    no pool
```

```
        drop-tail
        low
        percent-reduction-from-mbs default
    exit
exit
queue 12 auto-expedite multipoint create
    rate 100 cir 25 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
exit
exit
queue 13 auto-expedite multipoint create
    rate 100 cir 100 fir 0
    mbs 50
    cbs 10
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
exit
exit
queue 14 auto-expedite multipoint create
    rate 100 cir 100 fir 0
    mbs 50
    cbs 10
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
exit
exit
queue 15 auto-expedite multipoint create
    rate 100 cir 10 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
exit
exit
queue 16 auto-expedite multipoint create
    rate 100 cir 10 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
```

```
        exit
    exit
queue 17 auto-expedite multipoint create
    rate 100 cir 0 fir 0
    mbs 50
    cbs 1
    no pool
    drop-tail
        low
            percent-reduction-from-mbs default
    exit
    exit
exit
queue 18 auto-expedite multipoint create
    rate 100 cir 25 fir 0
    mbs 50
    cbs 3
    no pool
    drop-tail
        low
            percent-reduction-from-mbs default
    exit
    exit
exit
queue 19 auto-expedite multipoint create
    rate 100 cir 25 fir 0
    mbs 50
    cbs 10
    no pool
    drop-tail
        low
            percent-reduction-from-mbs default
    exit
    exit
exit
queue 20 auto-expedite multipoint create
    rate 100 cir 25 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
        low
            percent-reduction-from-mbs default
    exit
    exit
exit
queue 21 auto-expedite multipoint create
    rate 100 cir 100 fir 0
    mbs 50
    cbs 10
    no pool
    drop-tail
        low
            percent-reduction-from-mbs default
    exit
    exit
exit
queue 22 auto-expedite multipoint create
    rate 100 cir 100 fir 0
```

```
        mbs 50
        cbs 10
        no pool
        drop-tail
        low
            percent-reduction-from-mbs default
        exit
    exit
exit
queue 23 auto-expedite multipoint create
    rate 100 cir 10 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
        percent-reduction-from-mbs default
    exit
exit
exit
queue 24 auto-expedite multipoint create
    rate 100 cir 10 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
        percent-reduction-from-mbs default
    exit
exit
exit
queue 25 auto-expedite multipoint create
    rate 100 cir 0 fir 0
    mbs 50
    cbs 1
    no pool
    drop-tail
    low
        percent-reduction-from-mbs default
    exit
exit
exit
queue 26 auto-expedite multipoint create
    rate 100 cir 25 fir 0
    mbs 50
    cbs 3
    no pool
    drop-tail
    low
        percent-reduction-from-mbs default
    exit
exit
exit
queue 27 auto-expedite multipoint create
    rate 100 cir 25 fir 0
    mbs 50
    cbs 10
    no pool
    drop-tail
```

```
        low
        percent-reduction-from-mbs default
    exit
exit
queue 28 auto-expedite multipoint create
    rate 100 cir 25 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
exit
queue 29 auto-expedite multipoint create
    rate 100 cir 100 fir 0
    mbs 50
    cbs 10
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
exit
queue 30 auto-expedite multipoint create
    rate 100 cir 100 fir 0
    mbs 50
    cbs 10
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
exit
queue 31 auto-expedite multipoint create
    rate 100 cir 10 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
exit
queue 32 auto-expedite multipoint create
    rate 100 cir 10 fir 0
    mbs 25
    cbs 3
    no pool
    drop-tail
    low
    percent-reduction-from-mbs default
    exit
exit
```

```
exit
fc af create
    queue 3
    multicast-queue 11
    broadcast-queue 19
    unknown-queue 27
exit
fc be create
    queue 1
    multicast-queue 9
    broadcast-queue 17
    unknown-queue 25
exit
fc ef create
    queue 6
    multicast-queue 14
    broadcast-queue 22
    unknown-queue 30
exit
fc h1 create
    queue 7
    multicast-queue 15
    broadcast-queue 23
    unknown-queue 31
exit
fc h2 create
    queue 5
    multicast-queue 13
    broadcast-queue 21
    unknown-queue 29
exit
fc l1 create
    queue 4
    multicast-queue 12
    broadcast-queue 20
    unknown-queue 28
exit
fc l2 create
    queue 2
    multicast-queue 10
    broadcast-queue 18
    unknown-queue 26
exit
fc nc create
    queue 8
    multicast-queue 16
    broadcast-queue 24
    unknown-queue 32
exit
-----
*A:PE>config>qos>shared-queue#
```


10.4 Shared-Queue QoS Policy Command Reference

10.4.1 Command Hierarchies

10.4.1.1 Configuration Commands

```
config
  — qos
    — shared-queue policy-name [create]
    — no shared-queue policy-name
      — description description-string
      — no description
      — fc fc-name [create]
      — no fc
        — 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] [multipoint] [create]
      — no queue queue-id
        — cbs percent
        — no cbs
        — drop-tail
          — low
            — percent-reduction-from-mbs percent
            — no percent-reduction-from-mbs
          — mbs percent
          — no mbs
        — pool pool-name
        — no pool
        — rate percent [cir percent] [fir percent]
        — no rate
```

10.4.1.2 Show Commands

```
show
  — qos
    — shared-queue [shared-queue-policy-name] [detail]
```

10.4.2 Command Descriptions

10.4.2.1 Configuration Commands

10.4.2.1.1 Generic Commands

description

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

10.4.2.1.2 Shared Queue QoS Commands

shared-queue

Syntax	shared-queue <i>policy-name</i> [create] no shared-queue <i>policy-name</i>
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 **fc** *fc-name* [create]
no fc

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](#) for undefined forwarding class values.

Parameters *fc-name* — The case-sensitive, system-defined forwarding class name for which policy entries will be created.
Values be, l2, af, l1, h2, ef, h1, nc

broadcast-queue

Syntax **broadcast-queue** *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. When 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 this 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 specify an existing multipoint queue defined in the **config>qos>sap-ingress** context policer-output-queues profile. For the 7950 XRS, this is not configurable in the policer-output-queues profile.
Values 17 to 24

multicast-queue

Syntax	multicast-queue <i>queue-id</i>				
Context	config>qos>shared-queue>fc				
Description	This command configures the multicast forwarding type queue mapping for fc <i>fc-name</i>. The specified <i>queue-id</i> must exist within the policy as a multipoint queue before the mapping can be made. When the forwarding class mapping is executed, all multicast traffic on a SAP using this policy is forwarded using the <i>queue-id</i>. 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 this command sets the multicast forwarding type <i>queue-id</i> 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	<i>queue-id</i> — The <i>queue-id</i> parameter specified must be an existing, multipoint queue defined in the config>qos>sap-ingress context policer-output-queues profile. For the 7950 XRS, this is not configurable in the policer-output-queues profile. <table><tr><td>Values</td><td>9 to 16</td></tr><tr><td>Default</td><td>11</td></tr></table>	Values	9 to 16	Default	11
Values	9 to 16				
Default	11				

queue

Syntax	queue <i>queue-id</i> no queue
Context	config>qos>shared-queue>fc
Description	This command overrides the default unicast forwarding type queue mapping for fc <i>fc-name</i>. The specified <i>queue-id</i> must exist within the policy as a non-multipoint queue before the mapping can be made. When 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 <i>queue-id</i>. The no form of this command sets the unicast (point-to-point) <i>queue-id</i> back to the default queue for the forwarding class (queue 1).

Parameters	<i>queue-id</i> — The <i>queue-id</i> parameter specified must be an existing non-multipoint queue defined in the config>qos>sap-ingress context. For the 7950 XRS, this is not configurable in the policer-output-queues profile.
Values	Any valid non-multipoint <i>queue-id</i> in the policy including 1 and 3 through 32.
Default	1

unknown-queue

Syntax	unknown-queue <i>queue-id</i> no unknown-queue
Context	config>qos>shared-queue>fc
Description	This command configures the unknown unicast forwarding type queue mapping for fc <i>fc-name</i>. The specified <i>queue-id</i> must exist within the policy as a multipoint queue before the mapping can be made. When the forwarding class mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the <i>queue-id</i>. 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 <i>queue-id</i> back to the default of tracking the multicast forwarding type queue mapping.
Parameters	<i>queue-id</i> — The <i>queue-id</i> must be an existing, multipoint queue defined in the config>qos>sap-ingress context policer-output-queues profile. For the 7950 XRS, this is not configurable in the policer-output-queues profile. Values 25 to 32

queue

Syntax	queue <i>queue-id</i> [<i>queue-type</i>] [multipoint] [create] no queue <i>queue-id</i>
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 type is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or best effort 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 best effort forwarding classes are mapped to the queue (be, af, l1, or l2), the queue is treated as best effort by the hardware schedulers. The queue type can be overridden within the policy.

On the 7450 ESS, 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.

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 to 32

queue-type — The **expedite**, **best-effort**, and **auto-expedite** queue types are mutually exclusive. 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 forwarding class unicast traffic is mapped 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 preexisting multipoint queue to edit *queue-id* parameters.

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 XMA or 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 enqueueing packets. When 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. At egress, there are two additional RED slopes maintained in each buffer pool: the highplus slope that is used by inplus-profile packets, and the exceed slope, which is used by exceed-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 XMA or 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 being unused and should be avoided.

Default The cbs forwarding class defaults are listed in [Table 96](#).

Table 96 CBS Forwarding Class Default Values

Forwarding Class	Forwarding 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 Mbytes is the total buffer space in the buffer pool, a value of 10 would reserve 1 Mbytes (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 to 100

drop-tail

Syntax **drop-tail**

Context config>qos>shared-queue>queue

Description This command enters the context to configure queue drop tail parameters.

low

Syntax **low**

Context config>qos>shared-queue>queue>drop-tail

Description This command enters the context to configure the queue low drop-tail parameters. The low drop-tail defines the queue depth beyond which out-of-profile packets will not be accepted into the queue and will be discarded.

percent-reduction-from-mbs

Syntax **percent-reduction-from-mbs** *percent*
 no percent-reduction-from-mbs

Context config>qos>shared-queue>queue>drop-tail>low

Description This command configures the ingress shared queue low drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and **percent-reduction-from-mbs** is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value and will be discarded.

Default percent-reduction-from-mbs 10

Parameters *percent* — Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

mbs

Syntax	mbs <i>percent</i> 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 XMA or MDA forwarding class queue or egress network port forwarding class queue. The MBS value is used by a queue to determine whether it has exhausted its total allowed buffers while enqueueing packets. When 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 against 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 Table 97 .

Table 97 MBS Forwarding Class Default Values

Forwarding Class	Forwarding 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	<i>percent</i> — The percent of buffers from the total buffer pool space for the maximum amount of buffers, expressed as a decimal integer. If 10 Mbytes is the total buffers in the buffer pool, a value of 10 would limit the maximum queue size to 1 Mbyte (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 2 Mbytes.
Values	0 to 100

pool

Syntax	pool <i>pool-name</i> no pool
Context	config>qos>shared-queue>queue
Description	This command allows the queue to receive buffers from an explicit buffer pool instead of the default buffer pool. The specified <i>pool-name</i> 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. Note that named pools are not supported on the 7950 XSR, 7750 SRc4/12, 7750 SR-a4/a8, or 7750 SR-1e/2e/3e. If the specified pool-name does not exist on the XMA or MDA, the queue will be treated as pool orphaned and will be mapped to the appropriate default pool. When the pool is created on the MDA or port, the queue will be mapped to the new pool. The pool command may be used to either remove the queue from the pool, or to 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 in the queue command output using the pool keyword if it is configured. The no form of this command is used to remove a named pool association for the queue. When the pool name is removed, the queue will be placed on the appropriate default pool.
Parameters	<i>pool-name</i> — 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. The specified name must be an ASCII name string up to 16 characters.

rate

Syntax	rate <i>percent</i> [<i>cir percent</i>] [<i>fir percent</i>] no rate
Context	config>qos>shared-queue>queue
Description	This command defines the administrative PIR, the administrative CIR, and the administrative FIR parameters for the queue.

The PIR defines the percentage that the queue can transmit packets through the switch fabric (for 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 over-subscription factors or available bandwidth.

The CIR defines the percentage at which the system prioritizes the queue over other queues competing for the same bandwidth.

The **rate** command can be executed at any time, altering the PIR, CIR, and FIR for the queue created with the *queue-id*.

Parameters *percent* — Defines the percentage of the FP ingress capacity for the max rate allowed for the queue. When the **rate** command is executed, a valid percent (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 the value must be given as a positive integer.

Values 1 to 100, **max**

Default 100

cir percent — Defines the percentage of the FP ingress capacity for the CIR 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 the value must be given as a positive integer.

Values 0 to 100, **max**

Default 0

fir percent — Defines the percentage of the FP ingress capacity for the FIR allowed for the queue. When the **rate** command is executed, a FIR setting is optional. When the **rate** command has not been executed or the **fir** parameter is not explicitly specified, the default FIR (0) is assumed. Fractional values are not allowed and the value must be given as a positive integer. FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

Values 0 to 100

Default 0

10.4.2.2 Show Commands

The command outputs shown in this section are examples only; actual displays may differ depending on supported functionality and user configuration.

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** The following output is an example of shared-queue QoS policy information, and [Table 98](#) describes shared-queue QoS policy output fields.

Sample Output

```
*A:PE# show qos shared-queue

=====
Shared Queue Policies
=====
Policy-Id          Description
-----
default            Default Shared Queue Policy
policer-output-queues  Default Policer Output Shared Queue Policy
egress-pbr-ingress-queues  Egress PBR Ingress Shared Queue Policy
=====
*A:PE#
```

Table 98 **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.

11 Advanced QoS Policies

11.1 Overview

The `adv-config-policy` contains queue and policer configuration parameters that are not expected to be useful to most users. The policy only contains queue and policer child control parameters within a child-control node.

The parameters within the child-control node are intended to allow more precise control of the method that hierarchical virtual scheduling employs to emulate the effect of a scheduling context upon a member child queue or policer.

When 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).

While the system maintains default values for the advanced configuration parameters, no default `adv-config-policy` exists.

For information about the tasks and commands necessary to access the CLI and to configure and maintain the router, refer to CLI Usage chapter in the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Basic System Configuration Guide*.

11.2 Basic Configurations

A advanced QoS policy must conform to the following:

- Each advanced policy must have a unique policy ID.
- Default values can be modified but parameters cannot be deleted.

11.2.1 Creating an Advanced QoS Policy

Configuring and applying advanced policies is optional. If no advanced policy is explicitly applied to a SAP or IP interface, then no default advanced policy is applied.

To create a new advanced 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 child control parameters, which are intended to allow more precise control of the method that hierarchical virtual scheduling employs to emulate the effect of a scheduling context upon a member child queue or policer.

The following displays an example of an advanced policy configuration:

```
ALA-7>config>qos# info
#-----
echo "QoS Policies Configuration"
#-----
...
    adv-config-policy "childPolicy1" create
    description "Test"
    child-control
        offered-measurement
            add percent 30
            granularity percent 30
        exit
        bandwidth-distribution
            above-offered-cap percent 5
            granularity percent 5
        exit
    exit
exit
...
#-----
ALA-7>config>qos#
```

11.2.2 Applying Advanced Policies

Apply advanced policies to the following entities:

11.2.2.1 Queue Group

Use the following CLI syntax to apply advanced policies to ingress queue groups.

CLI Syntax: `config> qos queue-group-templates ingress queue-group
group1 policer 1 adv-config-policy child1`

Use the following CLI syntax to apply advanced policies to egress queue groups.

CLI Syntax: `config> qos queue-group-templates egress queue-group
group1 policer 1 adv-config-policy child1`

11.2.2.2 SAP Ingress

Use the following CLI syntax to apply advanced policies to an ingress SAP.

CLI Syntax: `qos sap-ingress 11 policer 1 adv-config-policy child1
queue 1 adv-config-policy child1`

11.2.2.3 SAP Egress

Use the following CLI syntax to apply advanced policies to an egress SAP.

CLI Syntax: `qos sap-egress 11 policer 1 adv-config-policy child1
queue 1 adv-config-policy child1`

11.2.3 Default Advanced Policy Values

The default policies cannot be edited or deleted. [Table 99](#) lists default advanced policy parameters.

Table 99 Advanced Policy Parameter Defaults

Field		Default
offered-measurement		
	high-rate-hold-time	0
	time-average-factor	0
	sample-interval	1

11.3 Service Management Tasks

This section discusses advanced QoS policy service management tasks.

11.3.1 Deleting QoS Policies

Delete advanced policies from the following entities.

11.3.1.1 Queue Group

Use the following CLI syntax to delete advanced policies from ingress queue groups.

CLI Syntax: `config> qos queue-group-templates ingress queue-group
group1 policer 1 no adv-config-policy`

Use the following CLI syntax to delete advanced policies from egress queue groups.

CLI Syntax: `config> qos queue-group-templates egress queue-group
group1 policer 1 no adv-config-policy`

11.3.1.2 SAP Ingress

Use the following CLI syntax to delete advanced policies from an ingress SAP.

CLI Syntax: `qos sap-ingress 11 policer 1 adv-config-policy child1
queue 1 no adv-config-policy`

11.3.1.3 SAP Egress

Use the following CLI syntax to delete advanced policies from an egress SAP.

CLI Syntax: `qos sap-egress 11 policer 1 adv-config-policy child1
queue 1 no adv-config-policy`

11.3.2 Copying and Overwriting Advanced Policies

An existing advanced policy can be copied, renamed with a new policy ID value, or used to 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 {adv-config-policy} source-policy-id
dest-policy-id [overwrite]`

11.3.3 Editing Advanced Policies

Existing advanced policies and entries in the CLI or NMS can be edited. 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, then write over the original policy.

11.4 Advanced QoS Policy Command Reference

11.4.1 Command Hierarchies

11.4.1.1 Configuration Commands

```

config
  — qos
    — adv-config-policy policy-name [create]
    — no adv-config-policy policy-name
      — child-control
        — bandwidth-distribution
          — above-offered-cap {percent percent-of-admin-pir | rate rate-in-kilobits-per-second}
          — no above-offered-cap
          — [no] enqueue-on-pir-zero
          — granularity {percent percent-of-admin-pir | rate rate-in-kilobits-per-second}
          — no granularity
          — internal-scheduler-weight-mode {default | force-equal | offered-load | capped-offered-load}
          — no internal-scheduler-weight-mode
          — [no] limit-pir-zero-drain
          — [no] lub-init-min-pir
        — offered-measurement
          — add percent percentage [min-only] [active-min-only]
          — add rate rate [min-only] [active-min-only]
          — no add
          — [no] fast-start
          — [no] fast-stop
          — granularity {percent percent-of-admin-pir | rate rate-in-kilobits-per-second}
          — no granularity
          — high-rate-hold-time seconds [active-min-only]
          — no high-rate-hold-time
          — max-decrement {percent percent-of-admin-pir | rate rate-in-kilobits-per-second}
          — no max-decrement
          — sample-interval sample-period
          — no sample-interval
          — time-average-factor taf-value [dec-only]
          — no time-average-factor
      — description description-string
      — no description

```

11.4.1.2 Operational Commands

```

config
— qos
    — copy adv-config-policy src-name dst-name [overwrite]

```

11.4.1.3 Show Commands

```

show
— qos
    — adv-config-policy [policy-name] [association]

```

11.4.2 Command Descriptions

11.4.2.1 Configuration Commands

11.4.2.1.1 Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>adv-config-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, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

11.4.2.2 Advanced Policy QoS Commands

adv-config-policy

Syntax **adv-config-policy** *policy-name* [**create**]
no adv-config-policy *policy-name*

Context config>qos

Description This command enters the context to configure an advanced QoS policy. This command contains only queue and policer child control parameters within a child-control node.

The parameters within the **child-control** node are intended to allow more precise control of the method that hierarchical virtual scheduling employs to emulate the effect of a scheduling context upon a member child queue or policer.

When 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).

While the system maintains default values for the advanced configuration parameters, no default **adv-config-policy** exists.

The **no** form of this command removes the specified advanced policy.

Default None

Parameters *policy-name* — The name of the advanced QoS policy. A policy-name must be specified and conform to the policy naming guidelines. If the specified name does not exist, the optional **create** keyword requirements are met and the total number of policies per system will not be exceeded, an **adv-config-policy** of that name will be created. If the specified name does exist, the system will switch context to that **adv-config-policy** for the purpose of modification of the policy's contents.

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

child-control

Syntax	child-control
Context	config>qos>adv-config-policy
Description	This command contains parameters that are intended to allow more precise control of the method that hierarchical virtual scheduling employs to emulate the effect of a scheduling context upon a member child queue or policer. This command edits the parameters that control the child requested bandwidth and parental bandwidth distribution for all policers and queues associated with the policy.

bandwidth-distribution

Syntax	bandwidth-distribution
Context	config>qos>adv-config-policy>child-control
Description	This command modifies or controls the bandwidth distribution phase of the parent virtual scheduler. This command edits the parameters that control the child given bandwidth for all policers and queues associated with the policy.

above-offered-cap

Syntax	above-offered-cap { <i>percent percent-of-admin-pir</i> <i>rate rate-in-kilobits-per-second</i> } no above-offered-cap
Context	config>qos>adv-config-policy>child-control>bandwidth-distribution
Description	This command is used to limit the operationally configured shaping or policing rate on the child associated with the policy. After the parent virtual scheduler or policer control policy determines the appropriate rate for a given child, a separate operation decides the actual PIR that should be configured for that child. When the parent determines that the distributed rate is equal to or less than the child's offered rate, the configured operational PIR will be equal to that determined rate. But when the parent determines that the child's offered rate is less than the available bandwidth the child could consume, the operational PIR may be set to a value larger than the distributed bandwidth. This extra rate is not currently used by the child since the offered rate is less. The system provides this extra bandwidth in case the child's offered rate increases before the next sampling interval is complete, in order to mitigate the periodic nature of the child's operational PIR adjustments. The increase in the offered rate is not subtracted from the parent's remaining distribution bandwidth for lower priority children, only the determined rate is considered consumed by the parent virtual scheduler or policer control policy instance. The actual operationally configured PIR will never be greater than the child's administratively defined PIR.

This 'fair share' PIR configuration behavior may result in the sum of the children's PIRs exceeding the aggregate rate of the parent. If this behavior violates the downstream QoS requirements, the **above-offered-cap** command may be used to minimize or eliminate the increase in the child's configured PIR.

If the **above-offered-cap** command is used with a percent-based value, the increase is a function of the configured PIR value on the policer or queue. In this case, care should be taken that the child is either configured with an explicit PIR rate (other than max) or the child's administrative PIR is defined using the percent-rate command with the local parameter enabled if an explicit value is not desired. When a maximum PIR is in use on the child, the system attempts to interpret the maximum child forwarding rate. This rate could be very large if the child is associated with multiple ingress or egress ports.

If the child's administrative PIR is modified while a percent based above-offered-cap is in effect, the system automatically uses the new relative limit value the next time the child's operational PIR is distributed.

When this command is not specified or removed, the child's operational 'fair share' operational PIR may be configured up to the child's administrative PIR, based on the actual parental bandwidth available at the child's priority level.

The **no** form of this command is used to remove a fair share operational PIR rate increase limit from all child policers and queues associated with the policy.

Parameters

percent-of-admin-pir — When the percent qualifier is used, the following percent-of-admin-pir parameter specifies the percentage of the child's administrative PIR that is used as the fair share increase limit. The new operational PIR result is capped by the child's PIR. If a value of 0 or 0.00 is used, the system will disable the fair share increase function and only configure the actual distribution rate. If a value of 100 or 100.00 is used, the system will interpret this equivalent to executing the **no above-offered-cap** command and return the fair-share operation to the default behavior.

Values 0.00 to 100.00

rate-in-kilobits-per-second — When the rate qualifier is used, the rate-in-kilobits-per-second parameter specifies an explicit rate, in kb/s, that are used as the limit to the child's fair share increase to the operational PIR. The new operational PIR result is capped by the child's PIR. If a value of 0 is used, the system will disable the fair share increase function and only configure the actual distribution rate.

Values 0 to 100,000,000

enqueue-on-pir-zero

Syntax [no] enqueue-on-pir-zero

Context config>qos>adv-config-policy>child-control>bandwidth-distribution

Description This command is used to enable queuing of new packets when HQoS determines that a queue should stop forwarding (operational PIR set to zero). The default behavior is to allow the queue to continue to use the previously determined operational PIR and set the queue's MBS (Maximum Burst Size) to zero. This prevents new packets from being admitted to the queue until the PIR zero case terminates. The new behavior when **enqueue-on-pir-zero** is enabled is to set the operational PIR to zero and leave the queue's MBS set to the normal value.

This command overrides the **limit-pir-zero-drain** command.

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

granularity

Syntax **granularity** {**percent** *percent-of-admin-pir* | **rate** *rate-in-kilobits-per-second*}
no granularity

Context config>qos>adv-config-policy>child-control>bandwidth-distribution

Description This command is used to create a step-like behavior where the operational PIR will round up to the nearest increment of the specified granularity before being applied to the child. The only exception is when the distributed bandwidth is less than 1% above a lower step value, in which case the lower step value is used.

This step-like behavior may be useful when the bandwidth used by an active child is well known. While the **above-offered-cap** command automatically adds a specified amount to the operational PIR of a child, the **granularity** command only increments the operational PIR to the next step value. While not expected to be used in conjunction, the **above-offered-cap** and **granularity** commands may be used simultaneously, in which case the above-offered-cap increase will be applied first, followed by the granularity rounding to the next step value.

If the **granularity** command is used with a percent-based value, the rounding up function of the configured PIR value on the policer or queue is based on the child's administrative PIR. In this case, care should be taken that the child is either configured with an explicit PIR rate (other than max) or the child's administrative PIR is defined using the percent-rate command with the local parameter enabled if an explicit value is not desired. When a maximum PIR is in use on the child, the system attempts to interpret the maximum child forwarding rate. This rate could be very large if the child is associated with multiple ingress or egress ports.

If the child's administrative PIR is modified while a percent based granularity is in effect, the system automatically uses the new relative rounding value the next time the child's operational PIR is determined.

When this command is not specified or removed, the system makes no attempt to round up the child's determined operational PIR.

The **no** form of this command is used to remove the operational PIR rounding behavior from all child policers and queues associated with the policy.

Parameters	<i>percent-of-admin-pir</i> — When the percent qualifier is used, the following percent-of-admin-pir parameter specifies the percentage of the child's administrative PIR that should be used as the rounding step value. If a value of 0 or 0.00 is used, the system will interpret this equivalent to no granularity. Values 0.00 to 100.00 <i>rate-in-kilobits-per-second</i> — When the rate qualifier is used, the following rate-in-kilobits-per-second parameter specifies an explicit rate, in kb/s, that should be used as the child's rounding step value. If a rate step of 0 is specified, the system interprets this equivalent to no granularity. Values 0 to 100,000,000

internal-scheduler-weight-mode

Syntax	internal-scheduler-weight-mode {default force-equal offered-load capped-offered-load} no internal-scheduler-weight-mode
Context	config>qos>adv-config>policy>child-control>bandwidth-distribution
Description	This command overwrites the internal scheduler weight configured on a card level.
Default	internal-scheduler-weight-mode default
Parameters	default — Specifies that queues are equally weighted, except at ingress for mixed-speed LAGs without per-fp-inq-queuing enabled and at egress for all mixed-speed LAGs, in which cases the queues are weighted based on port speed. force-equal — Specifies that the queues are always equally weighted. offered-load — Specifies that the queues are weighted based on observed offered load. capped-offered-load — Specifies that the queues are weighted based on observed offered load capped. by PIR

limit-pir-zero-drain

Syntax	[no] limit-pir-zero-drain
Context	config>qos>adv-config>policy>child-control>bandwidth-distribution
Description	This command is used to configure the system to use the minimum configurable PIR instead of an HQoS derived zero operational PIR. The default behavior is to allow the operational PIR of the queue to remain the last configured value while setting the queue MBS to zero (preventing queuing of newly arriving packets). Retaining the previous PIR value may cause a momentary burst above an aggregate rate associated with the queue as it drains. Using the limit-pir-zero-drain command causes the queue to drain at the lowest rate possible (typically 1 kb/s) that limits overrun situations.

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

lub-init-min-pir

Syntax	[no] lub-init-min-pir
Context	config>qos>adv-config-policy>child-control>bandwidth-distribution
Description	This command is used to initialize new queues associated with a LUB context to use a minimum PIR similar to the effect of the limit-pir-zero-drain command. When a queue is initially created in a LUB context, it defaults to a zero value PIR until HQoS has an opportunity to configure an offered rate based operational PIR. Enabling this command forces a minimum rate operational PIR to be applied to the queue for use by enqueued packets prior to an HQoS iteration.

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

offered-measurement

Syntax	offered-measurement
Context	config>qos>adv-config-policy>child-control
Description	This command modifies the offered rate measurement used to determine the bandwidth the queue or policer is requesting from its parent virtual scheduling context. This command modifies the parameters that control the child requested bandwidth for all policers and queues associated with the policy.

add

Syntax	add percent <i>percentage</i> [min-only] [active-min-only] add rate <i>rate</i> [min-only] [active-min-only] no add
Context	config>qos>adv-config-policy>child-control>offered-measurement
Description	This command is used to increase the measured rate of the policer or queue associated with the policy. The offered rate (capped by the administrative PIR configured on the queue or policer) is usually used unaltered by the parent virtual scheduler. The add command allows this measured rate to be increased by the specified amount or by a percentage of the administrative PIR. The resulting rate will not exceed the administrative PIR.

The parent scheduler uses the modified measured rate as the available work load for the queue or policer in determining how much bandwidth the child should receive from the bandwidth distribution algorithm.

One example of when an increase in the measured offered rate may be desired is when a queue or policer is handling VoIP traffic. A characteristic of VoIP is the step nature in how traffic is used. Each call typically adds a certain maximum amount to the overall load. By using the add command, the bandwidth required for the next added call may be included in the current measured rate. This allows the virtual scheduler to allocate sufficient bandwidth to the queue or policer so that when the call is made the scheduling algorithm does not need to run to increase the bandwidth.

A side effect of increasing the measured offered rate is that if the extra bandwidth is allocated by the virtual scheduler, the available bandwidth to lower priority queues or policers is diminished even though the extra allocated bandwidth may not be in use. If this is the case, the effect will be seen as an underrun in the aggregate output of the virtual scheduler.

If the add command is used with a percent based value, the increase is a function of the configured PIR value on the policer or queue. In this case, care should be taken that the child is either configured with an explicit PIR rate (other than max) or the child's administrative PIR is defined using the percent-rate command with the local parameter enabled if an explicit value is not desired. When a maximum PIR is in use on the child, the system attempts to interpret the maximum child forwarding rate. This rate could be very large if the child is associated with multiple ingress or egress ports.

Except for the overall cap on the offered input into the virtual scheduler, the child's administrative PIR has no effect on the calculated increase if an explicit rate is specified.

If the child's administrative PIR is modified while a percent based add is in effect, the system automatically uses the new relative increase value the next time the child's offered rate is determined.

When the add command is not specified or removed, the child's offered rate used by the child's virtual scheduler is not increased.

The **no** form of this command is used to remove an offered rate increase from all child policers and queues associated with the policy.

Parameters *percent-of-admin-pir* — When the percent qualifier is used, this parameter specifies the percentage of the child's administrative PIR that should be added to the child's offered rate. The new offered rate result is capped by the child's PIR. If a value of 0 or 0.00 is used, the system interprets this equivalent to no add.

Values 1.00 to 100.00

rate-in-kilobits-per-second — When the rate qualifier is used, this parameter specifies an explicit rate, in kb/s, that should be added to the child's offered rate. The new offered rate result is capped by the child's PIR. If a rate increase of 0 is specified, the system interprets this equivalent to no add.

Values 0 to 100,000,000

min-only — This optional parameter is used to reinterpret the increase as a minimum offered rate. When this option is enabled, the system uses the specified increase as a minimum offered rate even for inactive queues or policers associated with the policy.

active-min-only — When this optional parameter is specified, the respective rate or percentage is treated as the minimum offered rate for a queue only when the queue has an actual non-zero offered rate. This is intended to limit the artificial increase in offered rate to queues that are currently active. When a queue's measured offered rate drops to zero, the system stops enforcing the minimum value.

fast-start

Syntax	[no] fast-start
Context	config>qos>adv-config-policy>child-control>offered-measurement
Description	This command is used to enable fast detection of initial bandwidth on a child policer or queue associated with the policy. Multiple offered rate counter reads may be performed per the sampling interval. The system accumulates these counter values and evaluates the delta at the conclusion of the sampling interval. When fast-start is enabled, the system identifies all children associated with the policy that enter the inactive state (current offered rate is zero). Any inactive 'fast start' child that has a positive offered counter during a sampling period bypasses the normal sampling interval and does an immediate offered rate evaluation. This option is intended for use with children that would benefit from faster than normal startup detection, typically those of a real-time nature. When this parameter is not enabled, the system uses the normal sampling interval behavior of both newly active and currently active children. The no form of this command is used to restore the sampling-interval-based offered rate evaluation for newly active children.

fast-stop

Syntax	[no] fast-stop
Context	config>qos>adv-config-policy>child-control>offered-measurement
Description	This command is used to enable fast detection of lack of offered rate on a child policer or queue associated with the policy. Multiple offered rate counter reads may be performed per sampling interval. The system accumulates these counter values and evaluates the delta at the conclusion of the sampling interval. When fast-stop is enabled, the system bypasses the sampling interval for any currently active 'fast stop' child that has a zero offered counter measurement and does an immediate offered rate evaluation using the zero value. This option is intended for use with children where other children would benefit from faster than normal inactive detection, typically those of a real-time nature. When this parameter is not enabled, the system uses the normal sampling interval behavior of both newly inactive and currently active children.

The **no** form of this command is used to restore the sampling-interval-based offered rate evaluation for newly inactive children.

granularity

Syntax	granularity { percent <i>percent-of-admin-pir</i> rate <i>rate-in-kilobits-per-second</i> } no granularity
Context	config>qos>adv-config-policy>child-control>offered-measurement
Description	This command is used to adjust the sensitivity of the virtual scheduler to changes in the child offered rate. As the child offered rate is determined, it is compared to the previous offered rate. If the delta does not exceed the sensitivity threshold determined for the current offered rate, the change in offered rate is ignored for that iteration. While it is assumed that changing the offered rate change sensitivity will be a rare occurrence, it may be prudent to react to smaller changes in the offered rate of a particular child policer or queue. Another possible reason for changing the sensitivity is that it may be desired to lower the impact of changes in offered rate on the virtual scheduler for a particular child by raising the granularity. A side effect of higher sensitivity (lower granularity) is that the virtual scheduler may need to adjust the distributed bandwidth between all children more often, resulting in the possibility of lowering resources available to other virtual scheduler instances on the slot. A side effect of lower sensitivity (higher granularity) is that the parent virtual scheduler may distribute insufficient bandwidth to the child resulting in dropped packets. If the granularity command is used with a percent based value, the sensitivity is a function of the configured PIR value on the policer or queue. In this case, care should be taken that the child is either configured with an explicit PIR rate (other than max) or the child's administrative PIR is defined using the percent-rate command with the local parameter enabled if an explicit value is not desired. When a maximum PIR is in use on the child, the system attempts to interpret the maximum child forwarding rate. This rate could be very large if the child is associated with multiple ingress or egress ports. Except for the overall cap on the offered input into the virtual scheduler, the child's administrative PIR has no effect on the calculated sensitivity if an explicit rate is specified. If the child's administrative PIR is modified while a percent based granularity is in effect, the system automatically uses the new relative sensitivity value the next time the child's offered rate is determined. The no form of this command is used to restore the default offered rate sensitivity behavior to all child policers and queues associated with the policy.

Parameters	<i>percent-of-admin-pir</i> — When the percent qualifier is used, this parameter specifies the percentage of the child's administrative PIR that are used as the threshold sensitivity to offered rate change. If a value of 0 or 0.00 is used, the system will interpret this equivalent to no granularity. Values 1.00 to 100.00 <i>rate-in-kilobits-per-second</i> — When the rate qualifier is used, this parameter specifies an explicit rate, in kb/s, that are used as the child's offered rate change sensitivity value. If a rate sensitivity of 0 is specified, the system interprets this equivalent to no granularity. Values 0 to 100,000,000

high-rate-hold-time

Syntax	high-rate-hold-time <i>seconds</i> [active-min-only] no high-rate-hold-time
Context	config>qos>adv-config-policy>child-control>offered-measurement
Description	This command sets a time period that the current offered rate should be maintained for a child policer or queue when it is seen that the offered rate is decreasing. The offered measurement that triggers the hold time is used when the hold timer expires, unless a higher offered rate is seen in the interim. When a higher rate is observed, the hold timer is canceled and the higher offered rate is used immediately. A possible reason to define a hold timer for an offered rate is to allow a child queue is to dampen the effects of a child with a fluctuating rate on the virtual scheduler. This works similar to the max-decrement in that the child holds on to bandwidth from the virtual scheduler in case it may be needed in the near future. This parameter has no effect on an increase to the child's offered rate. If the rate increase is above the change sensitivity, the new offered rate is immediately used. When this command is not specified or removed, the virtual scheduler immediately reacts to measured decreases in offered load. The no form of this command is used to remove any currently configured hold time for all child policers and queues associated with the policy. When the hold time is removed, any current hold timers for child policers are automatically canceled.
Parameters	<i>seconds</i> — The hold time configured must be specified in seconds. A value of 0 is equivalent to no high-rate-hold-time. Default 0 Values 0 to 60

active-min-only — When this optional parameter is specified, the **high-rate-hold-time** command will accept the optional **active-min-only** parameter. Attempting to remove the active-min-only parameter from the **add** command, or removing the **add** command itself, will fail while **active-min-only** is enabled on the **high-rate-hold-time** command. When specified, the respective rate or percentage is treated as the minimum offered rate for a queue, only when the queue has an actual non-zero offered rate. This is intended to limit the artificial increase in offered rate to queues that are currently active. When a queue's measured offered rate drops to zero, the system stops enforcing the minimum value.

max-decrement

Syntax	max-decrement { <i>percent percent-of-admin-pir</i> <i>rate rate-in-kilobits-per-second</i> } no max-decrement
Context	config>qos>adv-config-policy>child-control>offered-measurement
Description	This command is used to limit how fast a child queue or policer can 'give up' bandwidth that it has been allotted from the virtual scheduler in a single iteration. If the child's new offered rate has decreased by more than the maximum decrement limit, the system ignores the new offered rate and instead uses the old offered rate less the maximum decrement limit. A possible reason to define a maximum decrement limit is to allow a child queue or policer to hold on to a portion of bandwidth that has been distributed by the parent virtual scheduler in case the child's offered rate fluctuates in an erratic manor. The max-decrement limit has a dampening effect to changes in the offered rate. A side effect of using a maximum decrement limit is that unused bandwidth allocated to the child queue or policer will not be given to another child as quickly. This may result in an underrun of the virtual scheduler's aggregate rate. The max-decrement limit has no effect on any increase in a child's offered rate. If the rate increase is above the change sensitivity, the new offered rate is immediately used. If the max-decrement command is used with a percent based value, the decrement limit will be a function of the configured PIR value on the policer or queue. In this case, care should be taken that the child is either configured with an explicit PIR rate (other than max) or the child's administrative PIR is defined using the percent-rate command with the local parameter enabled if an explicit value is not desired. When a maximum PIR is in use on the child, the system attempts to interpret the maximum child forwarding rate. This rate could be very large if the child is associated with multiple ingress or egress ports. Except for the overall cap on the offered input into the virtual scheduler, the child's administrative PIR has no effect on the calculated sensitivity if an explicit rate is specified. If the child's administrative PIR is modified while a percent based max-decrement is in effect, the system automatically uses the new relative maximum decrement limit value the next time the child's offered rate is determined.

When the max-decrement command is not specified or removed, the virtual scheduler does not limit a decreasing offered rate to a specific limit.

The **no** form of this command is used to remove any currently configured maximum decrement limit for all child policers and queues associated with the policy.

Parameters *percent-of-admin-pir* — When the percent qualifier is used, this parameter specifies the percentage of the child's administrative PIR that should be used as the decrement limit to offered rate change. If a value of 100 or 100.00 is used, the system will interpret this equivalent to **no max-decrement**.

Values 1.00 to 100.00

rate-in-kilobits-per-second — When the rate qualifier is used, this parameter specifies an explicit rate, in kb/s, that should be used as the child's offered rate change sensitivity value. If a rate sensitivity of 0 is specified, the system interprets this equivalent to **no max-decrement**.

Values 0 to 100,000,000

sample-interval

Syntax **sample-interval** *sample-period*
no sample-interval

Context config>qos>adv-config-policy>child-control>offered-measurement

Description This command is used to define the number of intervening sample periods before a new offered rate is measured and is only applicable when the policy is applied to a policer. By decreasing the sampling interval, the system will measure a child's new offered rate more frequently. Inversely, increasing the sampling interval causes the child's offered rate to be measured less frequently.

The overall number of offered rate measurements the system attempts within a given timeframe is not affected by the sample-interval command. If the system is asked to perform offered rate measurements more often on some policers, it will take longer to get to all children.

When this command is not specified or removed, the system evaluates the offered rate of each child after 1 sampling period.

The **no** form of this command is used to restore the sampling interval default of 1 sample period.

Parameters *sample-period* — The sample-periods parameter is specified as a whole number between 1 and 8. The value '1' represents the fastest sampling rate available and the value '8' represents the slowest sampling period available.

Default 1

Values 1 to 8

time-average-factor

Syntax	time-average-factor <i>taf-value</i> [dec-only] no time-average-factor
Context	config>qos>adv-config-policy>child-control>offered-measurement
Description	This command is used to weight the new offered rate with a portion of the previous offered rate. It would be expected that this command would mainly be used with the dec-only option enabled. The adjustment to the offered rate is performed using the following formula when taf-value is not set to '0': $\text{Adjusted\_Rate} = ((\text{Prev\_Offered\_Rate} \times (\text{taf-value} - 1)) + \text{New\_Offered\_Rate}) / \text{taf-value}$ If the dec-only option is specified, the adjustment is only applied when New_Offered_Rate is less than the Prev_Offered_Rate. When taf-value is set to '0', the adjustment is never applied. The no form of this command is used to remove the time average factor adjustments to new offered rate measurements.
Parameters	<i>taf-value</i> — The taf-value is specified as a whole number between 0 and 64. The value '0' has special meaning in that it disables the time average factor adjustment and has the same effect as no time-average-factor. Default 0 Values 0 to 64 dec-only — This keyword is an optional parameter. When enabled, the time average factor adjustment is only applied if the new offered rate is decreasing compared to the previous offered rate. If the new offered rate is greater than the previous offered rate, the adjustment is not applied.

11.4.2.3 Operational Commands

copy

Syntax	copy adv-config-policy <i>src-name dst-name</i> [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**
- adv-config-policy** — Indicates that the source policy ID and the destination policy ID are advanced policy IDs. Specify the source advanced policy ID that the copy command will attempt to copy from and specify the destination advanced policy ID to which the command will copy a duplicate of the policy.
 - overwrite** — Specifies that this policy is to replace the existing destination advanced policy. Everything in the existing destination policy will be overwritten with the contents of the source advanced policy. If **overwrite** is not specified, an error will occur if the destination policy ID exists, as shown here:

```
ALA-7>config>qos# copy adv-config-policy default spl
MINOR: CLI Destination "spl" exists - use {overwrite}
ALA-7>config>qos#overwrite
```

11.4.2.4 Show Commands

The command outputs in the following section are examples only; actual displays depend on supported functionality and user configuration.

adv-config-policy

- Syntax** **adv-config-policy** [*policy-name*] [**association**]
- Context** show>qos
- Description** This command displays advanced QoS policy information.
- Parameters**
 - policy-name* — The name of the advanced QoS policy.
 - association** — Displays associations related to the specified adv-config-policy.

- Output** The following output is an example of advanced QoS policy information, and [Table 100](#) describes the advanced QoS policy fields.

Sample Output

```
*A:PE# show qos adv-config-policy "adv1"
=====
QoS Advanced Configure Policy
=====
Policy Name           : adv1
Description           : (Not Specified)
OM Add PIR Prcnt      : 0.00           OM Add Rate           : 0
```

```

OM Add [min-only : false] [active-min-only : false]
OM Gran PIR Prcnt      : 0.00      OM Gran Rate      : 0
Max Dec PIR Prcnt      : 100.00     Max Dec Rate      : 0
Abv Off PIR Prcnt      : 100.00     Abv Off Rate      : max
BWD Gran PIR Prcnt      : 0.00      BWD Gran Rate      : 0
Fast Start             : Disabled    Fast Stop          : Disabled
Time Average Factor     : 0 [dec-only : false]
Hi Rate Hold Time       : 0 [active-min-only : false]
Sample Interval         : 1          Enqueue on PIR Zero : false
Limit PIR Zero Drain    : false      Init LUB with min PIR : false
Internal Sched Wght Mode: no-override
=====
*APE#

```

Table 100 Show QoS Scheduler-Hierarchy Customer Output Fields

Label	Description
Policy Name	Displays the policy name.
Description	Displays the policy description.
OM Add PIR Prcnt	Displays the percentage of the administrative PIR by which the measured rate is increased.
OM Add Rate	Displays the amount, in kb/s, by which the measured rate is increased.
OM Add	min-only — When this option is true, the system uses the specified increase as a minimum offered rate even for inactive queues or policers associated with the policy. active-min-only — When this option is true, the respective rate or percentage is treated as the minimum offered rate for a queue only when the queue has an actual non-zero offered rate. This is intended to limit the artificial increase in offered rate to queues that are currently active. When a queue's measured offered rate drops to zero, the system stops enforcing the minimum value.
OM Gran PIR Prcnt	Displays the percentage of the child's administrative PIR that is used as the threshold sensitivity to offered rate change. If a value of 0 or 0.00 is used, the system will interpret this equivalent to no granularity.
OM Gran Rate	Displays the explicit rate, in kb/s, that are used as the child's offered rate change sensitivity value. If a rate sensitivity of 0 is specified, the system interprets this equivalent to no granularity.
Max Dec PIR Prcnt	Displays the percentage of the child's administrative PIR that is used as the decrement limit to offered rate change. If a value of 100 or 100.00 is used, the system interprets this equivalent as no max-decrement .

Table 100 Show QoS Scheduler-Hierarchy Customer Output Fields

Label	Description (Continued)
Max Dec Rate	Displays the rate, in kb/s, that is used as the child's offered rate change sensitivity value. If a rate sensitivity of 0 is specified, the system interprets this equivalent to no granularity.
Abv Off PIR Prcnt	
Abv Off Rate	
BWD Gran PIR Prcnt	Displays the percentage of the child's administrative PIR that is used as the rounding step value. If a value of 0 or 0.00 is used, the system will interpret this equivalent to no granularity .
BWD Gran Rate	Displays the rate, in kb/s, that is used as the rounding step value. If a value of 0 is used, the system will interpret this equivalent to no granularity .
Fast Start	Displays whether fast detection of initial bandwidth on a child policer or queue associated with the policy is enabled.
Fast Stop	Displays whether fast detection of lack of offered rate on a child policer or queue associated with the policy is enabled.
Time Average Factor	Displays the time average factor used to calculate the weight of the new offered rate using a portion of the previous offered rate. If dec-only is true, the time average factor adjustment is only applied if the new offered rate is decreasing compared to the previous offered rate.
Hi Rate Hold Time	Displays the time period that the current offered rate is maintained for a child policer or queue when it is seen that the offered rate is decreasing. The offered measurement that triggers the hold time is used when the hold timer expires, unless a higher offered rate is seen in the interim. When a higher rate is observed, the hold timer is canceled and the higher offered rate is used immediately. If active-min-only is true, the respective rate or percentage is treated as the minimum offered rate for a queue, only when the queue has an actual non-zero offered rate.
Sample Interval	Displays the number of intervening sample periods before a new offered rate is measured.
Enqueue on PIR Zero	
Limit PIR Zero Drain	If true, the system to use the minimum configurable PIR instead of an HQoS derived zero operational PIR.

Table 100 **Show QoS Scheduler-Hierarchy Customer Output Fields**

Label	Description (Continued)
Init LUB with min PIR	If true, new queues associated with a LUB context are assigned a minimum rate operational PIR to be applied to the queue for use by enqueued packets prior to an HQoS iteration.
Internal Sched Wght Mode	Specifies the internal scheduler weight: default — Specifies that queues are equally weighted, except at ingress for mixed-speed LAGs without per-fp-inq-queuing enabled and at egress for all mixed-speed LAGs, in which cases the queues are weighted based on port speed. force-equal — Specifies that the queues are always equally weighted. offered-load — Specifies that the queues are weighted based on observed offered load. capped-offered-load — Specifies that the queues are weighted based on observed offered load capped. by PIR

12 QoS ATM Traffic Descriptor Profiles


12.1 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 Connections (PVC).

12.2 QoS Model 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 39](#).

Figure 39 Hierarchical Scheduling for ATM-Based Services



7750SR_001

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 use HQoS scheduler policy 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.

Any discard decision is 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 that include an ATM SAP (VLL, VPLS, IES, and VPRN).

12.3 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. UBR is a special case of UBR+MIR where MIR=0.

12.4 ATM Traffic Descriptors and QoS Parameters

Table 101 lists the supported ATM traffic descriptors.

Table 101 ATM Traffic Descriptors

Service Category	Traffic Descriptors
CBR	P0_1 PIR in kb/s (applies to CLP=0 and CLP=1 flows)
Rt-VBR and nrt-VBR	P0_1 and S0_1 PIR in kb/s (applies to CLP=0 & CLP=1 flows) SIR in kb/s (applies to CLP=0 & CLP=1 flows) MBS in cells (applies to CLP=0 and CLP=1 flows) P0_1 and S0 PIR in kb/s (applies to CLP=0 & CLP=1 flows) SIR in kb/s (applies to CLP=0 flow only) MBS in cells (applies to CLP=0 flow only)
UBR/UBR+MIR	P0_1 PIR in kb/s (applies to CLP=0 & CLP=1 flows) MIR in kb/s (applies to CLP=0 & CLP=1 flows)

12.5 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

12.6 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. An rt-VBR VC or an 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.

12.7 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.

12.8 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, when a cell is discarded by the ATM policing function, no other cells for that PDU are accepted, with exception of the “tail cell” that is sent to inform the far end that the end of a frame has arrived.

PPD is enabled on the ATM SAP part of an AAL5 SDU Apipe VLL service and on all services where the ATM SAP cell stream is reassembled, 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 IOM; packets that are reassembled in the ATM-capable MDAs take advantage of the application of the WRED congestion avoidance service queues associated with the ATM SAP.

Depending on the type of XMA or MDA that supports ATM encapsulation, other XMA- or MDA-specific packet congestion control mechanisms operating on per-SAP queues in the XMA or MDA are also applied.

12.9 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.

12.9.1 Creating 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#
```

12.9.2 Applying ATM-TD-Profile Policies

Apply ATM QoS traffic descriptor profiles to the following entities:

12.9.2.1 ATM VLL (Apipe) SAPs

Use the following CLI syntax to apply ATM QoS traffic descriptor profile policies to Apipe SAPs on ingress and egress. This command applies only to the 7750 SR and 7950 XRS.

CLI Syntax:

```
config>service>apipe>sap# atm
egress
    traffic-desc traffic-desc-profile-id
ingress
    traffic-desc traffic-desc-profile-id
```

12.9.2.2 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
```

12.9.2.3 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
```

12.9.2.4 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
```

12.9.2.5 VPRN SAPs

Use the following CLI syntax to apply ATM QoS traffic descriptor profile policies to VPRN SAPs on ingress and egress. This command applies only to the 7750 SR and 7950 XRS.

CLI Syntax:

```
config>service>vprn>if>sap# atm
egress
traffic-desc traffic-desc-profile-id
ingress
traffic-desc traffic-desc-profile-id
```

12.9.2.6 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
```

12.9.3 Default ATM-TD-Profile Policy Values

The default ATM QoS traffic descriptor profile is identified as “Default Traffic Descriptor”. The default profile cannot be edited or deleted. [Table 102](#) lists default profile parameters.

Table 102 ATM-TD-Profile Defaults

Field	Default
atm-td-profile <i>traffic-desc-profile-id</i>	1
description	“Default Traffic Descriptor”

Table 102 ATM-TD-Profile Defaults (Continued)

Field	Default
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#
```

12.10 Service Management Tasks

This section discusses QoS ATM traffic descriptor profile service management tasks.

12.10.1 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`

12.10.2 Copying and Overwriting Profiles

An existing profile can be copied, renamed with a new profile ID value, or used to 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#`

12.10.3 Editing QoS Policies

Existing policies and entries in the CLI can be edited. 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, then write over the original policy.

12.11 ATM QoS Policy Command Reference

12.11.1 Command Hierarchies

12.11.1.1 Configuration Commands

```
config
  — qos
    — atm-td-profile traffic-desc-profile-id [create]
    — 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
      — weight weight
      — no weight
```

12.11.1.2 Operational Commands

```
config
  — qos
    — copy atm-td-profile src-prof dst-prof [overwrite]
```

12.11.1.3 Show Commands

```
show
  — qos
    — atm-td-profile [traffic-desc-profile-id] [detail]
  — service
    — sap-using [ingress | egress] atm-td-profile td-profile-id
```

12.11.2 Command Descriptions

12.11.2.1 Configuration Commands

12.11.2.1.1 Generic Commands

description

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

12.11.2.1.2 ATM QoS Policy Commands

atm-td-profile

Syntax	atm-td-profile <i>traffic-desc-profile-id</i> [create] no atm-td-profile <i>traffic-desc-profile-id</i>
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 preconfigured 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, 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 this command deletes a given traffic profile. 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	atm-td-profile 1 — Default Traffic Descriptor (UBR, no traffic, no shaping)
Parameters	<i>traffic-desc-profile-id</i> — Index identifier for a traffic descriptor profile
Values	1 to 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 this command sets the CLP bit set to zero.
Default	no clp-tagging

descriptor-type

Syntax	descriptor-type <i>type</i>
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.
Default	Table 103 lists the default values of descriptor type based on a service category.

Table 103 Descriptor Type Default Values

Service Category	Default Descriptor Type
CBR	P0_1
UBR	P0_1
UBR with MIR	P0_1
rt-VBR or nrt-VBR	P0_1andS0_1

Parameters	<i>type</i> — defines interpretation of traffic parameters that are specified for this profile
	Values P0_1, P0_1andS0_Tag, P0_1andS0, P0_1andS0_1
	Table 104 describes these rules.

Table 104 Descriptor Type Rate Interpretations

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 the descriptor type to a value not compatible with the service category will generate an error.

policing

Syntax	[no] policing
Context	config>qos>atm-td-profile
Description	This command determines whether ingress traffic is policed. Policing is valid for CBR (PIR), RT-VBR (PIR and SIR), NRT-VBR, and UBR/UBR+MIR (PIR). This is cell-based policing.
Default	no policing

service-category

Syntax	service-category <i>service-category</i>
Context	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. The router supports the ATM service categories on ATM-capable MDAs listed in Table 105 .

Table 105 ATM Service Categories

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	service-category 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. [Table 106](#) lists which defaults apply for shaping depending on the service and UBR/UBR+MIR category.

Table 106 Default Shaping Values

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).

[Table 107](#) lists 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.

Table 107 Traffic Descriptor Parameter Support by Service Category

Service Category	SIR	PIR	MBS	MIR	CDVT
CBR	N/A	Y	N/A	N/A	Y
rt-VBR	Y	Y (must be $\geq$ SIR)	Y	N/A	Y
Nrt-VBR	Y	Y (must be $\geq$ SIR)	Y	N/A	Y
UBR	N/A	Y	N/A	N/A	N/A
UBR with MIR	N/A	Y (must be $\geq$ MIR)	N/A	Y (non-zero MIR specified)	N/A

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 kb/s, then weight = 1

If $32 \text{ kb/s} < \text{traffic rate} < 8160 \text{ kb/s}$, then weight = floor (traffic rate / 32)

If traffic rate ≥ 8160 kb/s, then weight = 255

The configuration of weight unit (32 kb/s) 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 this command restores traffic parameters to their defaults for a given service category.

ATM traffic parameter default values are listed in [Table 108](#) in kb/s.

Table 108 ATM Traffic Parameter Defaults

Service Category	Traffic Parameter Defaults
CBR:	
PIR	0
rt-VBR and nrt-VBR	
PIR	0
SCR	0
MBS	32
UBR (by default UBR is without MIR)	
PIR	0
MIR	0

- Parameters**
- sir** *sir-val* — Sustained Information Rate (including cell overhead) in kilobits per second.
Values 0 to 4294967295
- pir** *pir-val* — Peak Information Rate (including cell overhead) in kilobits per second.
Values 0 to 4294967295
- mir** *mir-val* — Minimum Desired Information Rate (including cell overhead) in kilobits per second.
Values 0 to 4294967295
- mbs** *mbs-val* — Maximum Burst Size in cells
Values 0 to 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 to 4294967295

weight

Syntax	weight <i>weight</i> no weight				
Context	config>qos>atm-td-profile				
Description	VCs within the VP tunnel are serviced by a single scheduler assigned to each VP tunnel. VCs within the shaped VP tunnel will be degraded from the originally-assigned service category to a common UBR service category (default traffic descriptor). Scheduling between VCs will be based on WRR with a weight parameter that can be explicitly configured in the ATM traffic descriptor profile. If weight is not specifically configured, the defaults are used. The explicitly configured weight parameter is used only on the ATM MDA in the <i>max16k-vc</i> mode. On all other ATM capable MDAs (ASAP or ATM MDA in <i>max8k-vc</i> mode), the weight parameter is ignored.				
Parameters	weight — Specifies the relative weight of the VC. <table> <tr> <td>Values</td><td>1 to 255</td></tr> <tr> <td>Default</td><td> VC degraded from CBR = weight 10 VC degraded from rt-VBR = weight 7 VC degraded from nrt-VBR = weight 5 VC degraded from UBR+ = weight 2 VC degraded from UBR = weight 1 </td></tr> </table>	Values	1 to 255	Default	VC degraded from CBR = weight 10 VC degraded from rt-VBR = weight 7 VC degraded from nrt-VBR = weight 5 VC degraded from UBR+ = weight 2 VC degraded from UBR = weight 1
Values	1 to 255				
Default	VC degraded from CBR = weight 10 VC degraded from rt-VBR = weight 7 VC degraded from nrt-VBR = weight 5 VC degraded from UBR+ = weight 2 VC degraded from UBR = weight 1				

12.11.2.1.3 Operational Commands

copy

Syntax	copy atm-td-profile <i>src-prof dst-prof</i> [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 to 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#
```

12.11.2.1.4 Show Commands

The command outputs in the following section are examples only; actual displays depend on supported functionality and user configuration.

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 to 1000

detail — Displays detailed policy information including policy associations.

Output The following output is an example of ATM traffic descriptor profile information, and [Table 109](#) describes ATM traffic descriptor profile show command output.

Sample

```
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#

```

Table 109 ATM Traffic Descriptor Profile Command Output Fields

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 kb/s.
PIR	Displays the peak cell rate in kb/s.
MIR	Displays the Minimum Desired Cell Rate in kb/s.
MBS	Displays the maximum burst size in cells.
Shaping	Displays whether shaping is enabled or disabled for the traffic descriptor profile.

Table 109 ATM Traffic Descriptor Profile Command Output Fields

Label	Description
Entities using TDP-ID	Displays the number of entities using the ATM traffic descriptor.
None	Indicates that the parameter is not applicable for the configured service category.

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 to 65535

filter *filter-id* — The ingress or egress Filter Policy ID for which to display matching SAPs.

Values 1 to 65535

sap-id — Specifies the physical port identifier portion of the SAP definition.

Values The following values apply to the 7750 SR and 7950 XRS:

<i>sap-id</i> :	null	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i> <i>aps-id</i>]
	dot1q	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i> <i>aps-id</i>]: <i>qtag1</i>
	qinq	[<i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>
	atm	[<i>port-id</i> <i>aps-id</i>]: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
	frame	[<i>port-id</i> <i>aps-id</i>]: <i>dlci</i>
	cisco-hdlc	<i>slot/mda/port.channel</i>
	cem	<i>slot/mda/port.channel</i>
	ima-grp	<i>bundle-id</i> [: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>]
	port-id	<i>slot/mda/port</i> [. <i>channel</i>]

bundle-id	bundle-type-slot/mda.bundle-num
	bundle keyword
	type ima, ppp
	bundle-num 1 to 256
bpgrp-id	bpgrp-type-bpgrp-num
	bpgrp keyword
	type ima, ppp
	bpgrp-num 1 to 1280
aps-id	aps-group-id[.channel]
	aps keyword
	group-id 1 to 64
ccag-id	ccag-id.path-id[cc-type]:cc-id
	ccag keyword
	id 1 to 8
	path-id a, b
	cc-type .sap-net, .net-sap
	cc-id 0 to 4094
lag-id	lag-id
	lag keyword
	id 1 to 800
qtag1	0 to 4094
qtag2	*, 0 to 4094
vpi	NNI: 0 to 4095
	UNI: 0 to 255
vci	1, 2, 5 to 65535
dlci	16 to 1022

The following values apply to the 7450 ESS:

Values

sap-id	null	[port-id lag-id]
	dot1q	[port-id lag-id]:qtag1
	qinq	[port-id lag-id]:qtag1.qtag2
	frame	[port-id]:dlci
	port-id	slot/mda/port[.channel]
	ccag-id	ccag-id.path-id[cc-type]:cc-id
	ccag	keyword

	<i>id</i>	1 to 8
	<i>path-id</i>	a, b
	<i>cc-type</i>	.sap-net, .net-sap
	<i>cc-id</i>	0 to 4094
lag-id	<i>lag-id</i>	
	<i>lag</i>	keyword
	<i>id</i>	1 to 800
qtag1		0. to 4094
qtag2		*, 0 to 4094
dlci		16 to 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 to 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 The following output is an example of service SAP information, and [Table 110](#) describes show service SAP output fields.

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#
```

Table 110 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.

Table 110 **SAP Output Fields (Continued)**

Label	Description
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.

13 Named Pools

13.1 Overview

The named buffer pool feature allows for the creation of named buffer pools at the XMA, MDA, and port level. Named pools allow for a customized buffer allocation mode for ingress and egress queues that goes beyond the default pool behavior.

Named pools are defined within a named pool policy. The policy contains a q1-pools context that is used to define port allocation weights and named pools for buffer pools on Q1-based IOMs (all IOMs that are supported). The policy may be applied at either the port or XMA or MDA level at which time the pools defined within the policy are created on the port or XMA/MDA. When the policy is applied at the XMA or MDA level, XMA or MDA named pools are created. XMA or MDA named pools will typically be used when either a pool cannot be created per port or when the buffering needs of queues mapped to the pool are not affected by sharing the pool with queues from other ports. XMA or MDA named pools allow buffers to be efficiently shared between queues on different ports mapped to the same pool. However, with XMA or MDA named pools, very active queues on one port could deplete buffers in the pool, with queues on other ports experiencing buffer starvation. Port named pools are created when the policy is applied at the port level and allow for a more precise application of the buffer space allocated for a physical port. XMA or MDA pool names do not need to be unique. If a name overlap exists, the port pool will be used. The same pool name may be created on multiple ports on the same XMA or MDA.

The named pool policy is applied at the XMA or MDA ingress and egress level and at the ingress and egress port level. Each XMA or MDA within the system is associated with a forwarding plane traffic manager that has support for a maximum of 57 buffer pools. The following circumstances affect the number of named pools that can be created per XMA or MDA (these circumstances may be different between ingress and egress for the XMA or MDA):

- The forwarding plane can be associated with multiple XMAs or MDAs (each XMA or MDA has its own named pools).
- A single system-level pool for system-created queues is allocated.
- Each system must have default pools for queues that are not explicitly mapped or are incorrectly mapped to a named pool.
- Default pools for most IOM types (separate for ingress and egress).
- Access pool.
- Network pool.

- The number of named per-port pools is dependent on the number of ports the XMA or MDA supports, which is variable per XMA or MDA type.
- Per-port named pools cannot be used by ingress network queues, but pools defined in a named pool policy defined on an ingress all network port are still created.
 - Ingress network queues use the default network pool or XMA or MDA named pools.
 - Ingress port buffer space allocated to network mode ports is included in the buffers made available to ingress XMA or MDA named pools.
 - Ingress port buffer space on channelized ports associated with network bandwidth is included in the buffers made available to ingress XMA or MDA named pools
 - Ingress port named pools are only allocated buffers when the port is associated with some access mode bandwidth
- Per-port named pools on ports aggregated into a LAG are still created per physical port
- Default, named XMA or MDA and named per-port pools are allocated regardless of queue provisioning activity associated with the pool

If the named pool policy is applied to an XMA or MDA or port that cannot create every pool defined in the policy, the policy application attempt will fail. Any preexisting named pool policy on the XMA, MDA, or port will not be affected by the failed named pool policy association attempt.

When buffer pools are being created or deleted, individual queues may need to be moved to or from the default pools. When a queue is being moved, the traffic destined to the queue is first moved temporarily to a 'fail-over' queue. Then the queue is allowed to drain. When the queue is drained, the statistics for the queue are copied. The queue is then returned to the free queue list. A new queue is then created associated with the appropriate buffer pool, the saved stats are loaded to the queue then the traffic is moved from the fail-over queue to the new queue. While the traffic is being moved between the old queue to the fail-over queue then to the new queue, some out of order forwarding may be experienced. Also, any traffic forwarded through the fail-over queue will not be accounted for in billing or accounting statistics. A similar action is performed for queues that have the associated pool name added, changed or removed. This only applies to where fail-over queues are currently supported.

The first step in allowing named pools to be created for an XMA or MDA is to enable 'named-pool-mode' at the IOM level (config card slot-number named-pool-mode). Named pool mode may be enabled and disabled at any time. When MDAs are currently provisioned on the IOM, the IOM is reset to allow all existing pools to be deleted and the new default, named XMA or MDA, and named port pools to be

created and sized. If MDAs are not currently provisioned (as when the system is booting up), the IOM is not reset. When named pool mode is enabled, the system changes the way that default pools are created. The system no longer creates default pools per port, instead, a set of per forwarding plane level pools are created that are used by all queues that are not explicitly mapped to a named pool.

After the IOM has been placed into named pool mode, a named pool policy must be associated with the ingress and egress contexts of the XMA or MDA or individual ports on the XMA or MDA for named pools to be created. There are no named pools that exist by default.

Each time the default pool reserve, aggregate XMA or MDA pool limit or individual pool sizes is changed, buffer pool allocation must be reevaluated.

Pools may be deleted from the named pool policy at any time. Queues associated with removed or non-existent pools are mapped to one of the default pools based on whether the queue is access or ingress. The queue is flagged as pool-orphaned until either the pool comes into existence, or the pool name association is changed on the pool.

An ingress or egress port-managed buffer space is derived from the port's active bandwidth. Based on this bandwidth value compared to the other port's bandwidth value, the available buffer space is given to each port to manage. It may be desirable to artificially increase or decrease this bandwidth value to compensate for how many buffers are actually needed on each port. If one port has very few queues associated with it and another has many queues associated, the commands in the port's "modify-buffer-allocation-rate" CLI context may be used to move one port's bandwidth up, and another port's bandwidth down. As provisioning levels change between ports, the rate modification commands may be used to adapt the buffer allocations per port.

Buffer allocation rate modification is supported for both standard and named pool mode buffer allocation methods.

The system allocates buffers based on the following criteria:

- "named-pool-mode" setting on the IOM.
- Amount of path bandwidth on channelized ports.
- Existence of queues provisioned on the port or channel.
- Current speed of each port.
- Each ports "ing-percentage-of-rate" and "egr-percentage-of-rate" command setting.
- The port-allocation-weights setting for default, XMA or MDA, and port.
- The ports division between network and access bandwidth.

-
- Each individual named pool's network-allocation-weight and access-allocation-weight.

System reserved named pool names (cannot be used when configuring a named pool) are: **default**, **SAP Shared**, and **MC Path Mgmt**.

13.2 Named Pool Mode for IOM3-XP Card

IOM3-XP has one forwarding complex for MDA1 and MDA2. The total available buffer space is divided in the ingress and egress direction. The total ingress buffer space is shared by both MDAs in the ingress direction and the total egress buffer space is shared by both MDAs in the egress direction.

Each XMA or MDA can use a different named pool policy.

Network ingress queues can only use the MDA1 named pools. If named pools are configured for MDA2, they will not be used by network ingress queues. Network ingress queues configured to use MDA2 named pools will be considered pool orphaned. To check for orphan queues, use the command **show mda mda qos ingress orphaned-queues**. The same restriction applies for SAP shared queues using named pools.

Each named pool policy can have a maximum of 57 named pools configured.

The total number of named pools that can be configured per IOM3-XP is 245. The named pool usage per card can be checked with the **tools dump resource-usage card 1 all | match pools** command.

```
*A:PE# tools dump resource-usage card 1 all | match "Named Pools"
          Dynamic Named Pools +          245          0          245
    Dynamic Named Pools (in use by Ingress) -          0
    Dynamic Named Pools (in use by Egress) -          0
*A:PE#
```

13.3 Basic Configurations

In a basic named pool QoS policy, the default values can be modified but parameters cannot be deleted.

13.3.1 Creating a Named Pool QoS Policy

When a new named pool policy is created, the named pool policy ID value must be defined. The system does not dynamically assign a value.

13.3.1.1 Named Pool Configuration Procedure

Step 1: Configure the named pool policy.

CLI Syntax:

```
config# configure qos named-pool-policy 3pools create
q1-pools
    pool p1 create
    exit
    pool p2 create
    exit
    pool p3 create
    exit
exit all
```

Step 2: Apply the named pool policy on ingress and/or egress XMA or MDA and/or port.

Since the named pool mode is not yet active on the card, all queues are drawing buffer from the default pool.

Configure the queues to get buffers from a named pool.

Configure the named pool policy.

CLI Syntax:

```
config# network queue
configure qos
    copy network-queue default 15
    network-queue 15
    queue 1 pool p1
exit all
```

Step 3: Configure the queue profile to be used by the respective applications or port.

Configure the named pool policy.

CLI Syntax: `config# configure card 1 mda 2 network ingress queue-policy 15`

Step 4: Turn on the named pool on the card. The card will reboot and the named pool mode will be active on the card.

CLI Syntax: `configure card 1 named-pool-mode now`

To check the pools after named pool mode was enabled, use “show pools mda”. There are no port pools unless a port has configured a named pool policy. Only named pools that have active queues associated will be shown with a non-zero size; this means that the named pool was created. If a named pool is configured but has no active queue associated, the size of the pool is zero. The named pool would be instantiated but not created. An example is shown below.

Example:

```
A:PE# show pools 1/2
=====
```

Type	Id-FP	App.	Pool Name	Actual Admin	ResvCBS ResvCBS	PoolSize
MDA	1/2-1	Acc-Ing	default	138240		445440
				Sum		
MDA	1/2-1	Acc-Ing	MC Path Mgmt	96000		192000
				50%		
MDA	1/2-1	Acc-Egr	default	49920		161280
				Sum		
MDA	1/2-1	Net-Ing	default	0		0
				Sum		
MDA	1/2-1	Net-Egr	default	345600		798720
				Sum		
MDA	1/2-1	Ingress	p1	46080		149760
			Policy: 3pools	30%		
MDA	1/2-1	Ingress	p2	46080		149760
			Policy: 3pools	30%		
MDA	1/2-1	Ingress	p3	46080		149760
			Policy: 3pools	30%		

```
=====
A:PE#
```

13.3.2 Allocation Steps

Whether one or multiple MDAs share the same buffer space, the buffer space is portioned out on a per port basis. Each port gets an amount of buffering that is its fair share based on the port's bandwidth compared to the overall active bandwidth. This is identical to current behavior.

This mechanism takes the buffer space available and divides it into a portion for each port, based on the ports active bandwidth relative to the amount of active bandwidth for all ports associated with the buffer space. The number of ports sharing the same buffer space depends on the type of IOM the pools are being created on and the type of MDAs populated on the IOM. An active port is considered to be any port that has an active queue associated. When a queue is created for the port, the system will allocate the appropriate amount of buffer space to the port. This process is independently performed for both ingress and egress.

Normally, the amount of active bandwidth is considered as opposed to total potential bandwidth for the port when determining the ports fair share. If a port is channelized and not all bandwidth is allocated, only the bandwidth represented by the configured channels with queues configured is counted towards the bandwidth represented by the port. Also, if a port may operate at variable speeds (as in some Ethernet ports), only the current speed is considered. The number of buffers managed by a port may change due to queue creation and deletion, channel creation and deletion, and port speed variance on the local port or other ports sharing the same buffer space.

After the active bandwidth is calculated for the port, the result may be modified through the use of the 'ing-percentage-of-rate' and 'egr-percent-of-rate' commands. The default value of each is 100%, which allows the system to use all of the ports active bandwidth when deciding the relative amount of buffer space to allocate to the port. When the value is explicitly modified, the active bandwidth on the port is changed according to the specified percentage. If a value of 50% is given, the ports active bandwidth will be multiplied by .5; if a value of 150% is given, the active bandwidth will be multiplied by 1.5. This capability is independent of named pool mode. The ports rate percentage parameters may be modified at any time.

When named pool mode is configured on the buffer space, the ingress and egress chunk of buffering assigned to a port is now split into three smaller chunks for default pools, mda named pools, and port named pools. The way the buffering is split into the three smaller chunks is based on the 'port-allocation-weights' given in the named-pool-policy. The weights may come from either the XMA- or MDA-level applied named pool policy or the local port applied named pool policy. If a named pool policy is assigned on both locations, the defined 'port-allocation-weights' from the port associated policy will apply. Any of the weights may be set to '0', indicating that none of the buffers allocated to the port should be given to the pool category. If only XMA or MDA named pools are created, the port weight should be set to '0'; if

only port named pools are created, the XMA or MDA weight should be set to '0'. Setting the default weight to '0' should be done with care as any queues without named pool definitions or queues with non-existent pool names use the default pools. The weights are summed, then each individual weight is divided by the sum. The result is multiplied by the buffer space managed by port to determine the amount of buffer space given to each category.

The portion of buffering set aside for default pools is applied to the default network and access pools based solely on the amount of bandwidth in network and access modes on the physical port. Combining the default pool buffering chunks from all ports results in the aggregate size of the default pools. For Ethernet ports, all the bandwidth is either network or access (there is never both). For channelized ports, the amount given to network and access default pools is based on the ratio between access and network bandwidth and not based on weights defined within the pools. The default pool configuration is not changed when named pool mode is enabled.

The portion of buffering managed by the port relative to the ports XMA or MDA weight parameter is separated into two sub-portions based on the ratio between network and access bandwidth on port. If the port is all network or access bandwidth, that opposite sub-portion will be empty. The sub-portion associated with network bandwidth on the port and the sub-portion associated with access bandwidth on the port are individually summed with the network and access XMA or MDA named pool bandwidth set aside by the other ports on the XMA or MDA. The outcome is a total amount of access and a total amount of network bandwidth that will be given to XMA or MDA named pools based on each pools network and access weights.

The portion of buffering assigned to port named pools based on the port weight is also separated into two sub-portions in the same manner as step 4 in [Named Pool Configuration Procedure](#). The network and access bandwidth defines the ratio of bandwidth population between the network and access sub-portions. If the port is all network or access bandwidth, that opposite sub-portion will be empty. Each sub-portion is divided between the port named pools based on each pools network and access defined weights.

On the ingress side, network queues may only be associated with the default ingress network pool or one of the XMA or MDA named pools. Because ingress network queues may not use ingress port-based pools, the port's network sub-portion is added to the ports XMA or MDA named pool network sub-portion to be distributed to the named XMA or MDA pools based on each pools network weight.

13.4 Named Pools QoS Policy Command Reference

13.4.1 Command Hierarchies

13.4.1.1 Configuration Commands

```
config
  — qos
    — named-pool-policy policy-name [create]
    — no named-pool-policy policy-name
      — description description-string
      — no description
    — q1-pools
      — pool pool-name [create]
      — no pool pool-name
        — amber-alarm-threshold percentage
        — no amber-alarm-threshold
        — application-weights
          — access-allocation-weight weight
          — no access-allocation-weight
          — network-allocation-weight weight
          — no network-allocation-weight
        — description description-string
        — no description
        — red-alarm-threshold percentage
        — no red-alarm-threshold
        — resv-cbs percentage-of-pool-size [amber-alarm-action step
          percent max maximum]
        — no resv-cbs
        — slope-policy slope-policy-name
        — no slope-policy
      — port-allocation-weights default weight mda weight port weight
      — no port-allocation-weights
```

13.4.1.2 Show Commands

```
show
  — qos
    — named-pool-policy
    — named-pool-policy pool-policy-name detail
    — named-pool-policy pool-policy-name [association]
```

```

show
— card card-slot detail | match “Named Pool Mode”
— mda slot[/mda] [detail]
— pools port-id
— port port-id detail

```

13.4.2 Command Descriptions

13.4.2.1 Configuration Commands

13.4.2.1.1 Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>named-pool-policy config>qos>named-pool-policy>q1-pools>pool
Description	The description command is used to define an informational ASCII string. The string value may be defined or changed at any time. The no form of this command removes the explicit description string.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc), the entire string must be enclosed within double quotes.

13.4.2.2 Named Pool Policy Creation

named-pool-policy

Syntax	named-pool-policy <i>policy-name</i> create no named-pool-policy <i>policy-name</i>
Context	config>qos

Description This command creates a template that may be applied at the XMA, MDA, or port level to create named buffer pools. The policy may be applied at either the ingress or egress context for the port, XMA, or MDA. Policies applied on the MDA will take effect only after the named pool mode is set on the IOM. Setting the IOM named pool on will reboot the card.

Within the policy, named pools may be defined in the q1-pools context indicating that the provisioned pools will be used on Q1-based hardware. When the policy is associated at the XMA or MDA level, named pools defined in the policy allow queues from any port to be associated. When the policy is associated at the port level, the named pools created are only available to queues associated with the port. Each pool defined allows the slope-policy, resv-cbs, access-allocation-weight, and network-allocation-weight parameters to be configured for the pool. The policy also manages the port-allocation-weights used to divide the buffers managed by the port between named pools local to the port, named pools on the ports XMA or MDA, and the default pools. The allocation weights for a given port are derived in the following way (lowest to highest preference):

1. Port default allocation weights
Default: default 50, mda 50, port 50
2. XMA or MDA named pool policy port allocation weights
3. Port named pool policy port allocation weights
4. A named-pool-policy that is currently applied to an XMA or MDA or port may not be deleted. All associations between the policy and XMA or MDA must be removed prior to deleting the policy.

Pools in the policy may be added or removed at any time. If the policy is currently associated with an XMA, MDA, or port, the system will first check to ensure necessary resources exist on the port, XMA, or MDA, before allowing the pool creation within the policy to proceed. If the pool cannot be added, the pool pool-name command will fail. When a new pool is created, the system will scan all pool orphaned queues for queues associated with the new pool name wherever the policy is currently applied. (A queue with a defined pool name that does not exist is placed on its appropriate default pool until the pool comes into existence).

The **no** form of this command removes a specific named pool policy from the system. If the named pool policy is currently associated with an ingress or egress XMA, MDA, or port, the command will fail. If the named pool policy does not exist, the command has no effect and does not return an error.

Parameters *named-pool-policy-name* — The named-pool-policy-name is required. Each named pool policy must be uniquely named within the system. Names of up to 32 ASCII characters composed of printable, 7-bit ASCII characters are supported. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. A named pool policy must exist prior to applying the policy name to an XMA, MDA, or port.

There is a limit of 1024 policies per system.

Values Up to 32 character ASCII string

create — The create keyword is required if creating a new named pool policy 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 named pool policy already exists.

q1-pools

Syntax	q1-pools
Context	config>qos>named-pool-policy
Description	The q1-pools command is used to enter the configuration node for Q1-oriented named buffer pools. The named pool policy will support contexts for configuring pools of other types when other pool types exist.

pool

Syntax	pool <i>pool-name</i> [create] no pool <i>pool-name</i>
Context	config>qos>named-pool-policy>q1-pools
Description	The pool command is used to create a new or edit an existing named pool within the policy. A CLI node is created for the named pool that contains the slope-policy and resv-cbs commands. A named pool created within the q1-pools context may be used by queues on any physical port, XMA, or MDA where the policy is applied that has Q1-based buffer pools. When the policy is applied on an XMA, MDA, or port, creating a new pool will fail if a pool resource is not available for the port, XMA, or MDA.

When creating a pool, the defined name must be unique within the policy. No other named pool may share the same name.

When the pool is created, any queues currently on a default pool with a specified pool name the same as the new pool will be moved from the default pool to the new pool.

The **no** form of this command removes a specific named pool from the policy. If an instance of the named pool is currently associated with a created queue, the queue will be moved to the appropriate default pool. When the pool is deleted, the pool is removed from both the policy and any instance of the pool on an XMA, MDA, or port. The pool buffers are freed and may be available for other named pools.

Parameters *pool-name* — The pool-name parameter is required. Each named pool must be uniquely named within the policy. Names of up to 32 ASCII characters composed of printable, 7-bit ASCII characters are supported. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. A named pool must be defined prior to creating a queue associated with the named pool.

Up to 57 named pools may be created per policy.

Values A string, up to 32 characters in length.

create — The create keyword is required if creating a new named pool within the policy when the system is configured, and requires 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 named pool already exists within the policy.

amber-alarm-threshold

Syntax **amber-alarm-threshold** *percentage*
no amber-alarm-threshold

Context config>qos>named-pool-policy>q1-pools>pool

Description This command configures the threshold for the amber alarm on an allowed over-subscription.

Users can selectively enable amber or red alarm thresholds. If both are enabled (nonzero) the red alarm threshold must be greater than the amber alarm threshold.

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

Default no amber-alarm-threshold

Parameters *percentage* — Specifies the amber alarm threshold.

Values 1 to 1000

application-weights

Syntax **application-weights**

Context config>qos>named-pool-policy>q1-pools>pool

Description The application-weights CLI node context contains the network and access allocation weights. The network and access application weights are used to divide the network and access buffer space available to the pools between each named pool. When the policy is applied at the XMA or MDA level, the network and access application weights are applied to the network and access buffer space given to the XMA or MDA named pools by the ingress or egress ports. When the policy is applied at the port level, the network and access application weights are applied to the local port network and access buffer space.

access-allocation-weight

Syntax **access-allocation-weight** *buffer-allocation-weight*
no access-allocation-weight

Context config>qos>named-pool-policy>q1-pools>pool>application-weights

Description The access-allocation-weight command is used to define the weight used when dividing access associated buffer space between the named pools. When the named pool is created on an XMA or MDA, the access associated buffer space is summed from all ports. The pool's access allocation weight is divided by the total access allocation weights from all named pools on the XMA or MDA. The resulting factor is multiplied by the summed port access associated buffer space to derive the amount of access buffers applied to the pool. When the named pool is created on a port, the weight is applied against the local ports access associated buffer space to derive the access buffers applied to the pool. A similar process is used with the network-allocation-weight. The total buffers applied to the pool are the sum of the access and network buffers given to the pool.

Changing the weight does not change the total buffers allocated to the pools, just the ratio of distribution between the pools.

A weight of '0' indicates that the pool will not receive any access associated buffers. If all pools on the port, XMA, or MDA have a access-allocation-weight equal to 0, the access associated buffer will not be used at that level.

The **no** form of this command is used to return the pools access allocation weight to the default value.

Default access-allocation-weight 50

Parameters *buffer-allocation-weight* — The buffer-allocation-weight parameter is required when executing the access-allocation-weight command. A value of 0 to 100 is accepted. The default weight is 50. The weight value may be changed at any time resulting in a redistribution of access associated buffers among the pools at the XMA, MDA, or port level.

Values 0 to 100

Default 50

network-allocation-weight

Syntax	network-allocation-weight <i>buffer-allocation-weight</i> no network-allocation-weight
Context	config>qos>named-pool-policy>q1-pools>pool>application-weights
Description	The network-allocation-weight command is used to define the weight used when dividing network associated buffer space between the named pools. When the named pool is created on an XMA or MDA, the network associated buffer space is summed from all ports. The pool's network allocation weight is divided by the total network allocation weights from all named pools on the XMA or MDA. The resulting factor is multiplied by the summed port network associated buffer space to derive the amount of network buffers applied to the pool. When the named pool is created on a port, the weight is applied against the local ports network associated buffer space to derive the network buffers applied to the pool. A similar process is used with the access-allocation-weight. The total buffers applied to the pool are the sum of the access and network buffers given to the pool. Changing the weight does not change the total buffers allocated to the pools, just the ratio of distribution between the pools. A weight of '0' indicates that the pool will not receive any network associated buffers. If all pools on the port, XMA, or MDA have a network-allocation-weight equal to 0, the network associated buffer will not be used at that level. The no form of this command returns the pools network allocation weight to the default value.
Default	network-allocation-weight 50
Parameters	<i>buffer-allocation-weight</i> — The buffer-allocation-weight parameter is required when executing the network-allocation-weight command. An integer value of 0 to 100 is accepted. The default weight is 50. The weight value may be changed at any time resulting in a redistribution of network associated buffers among the pools at the XMA, MDA, or port level. Default 50

red-alarm-threshold

Syntax	red-alarm-threshold <i>percentage</i> no red-alarm-threshold
Context	config>qos>named-pool-policy>q1-pools>pool
Description	This command configures the threshold for the red alarm on an allowed over-subscription. Users can selectively enable amber or red alarm thresholds. If both are enabled (nonzero) the red alarm threshold must be greater than the amber alarm threshold. The no form of this command reverts to the default value.

Default	no red-alarm-threshold
Parameters	<i>percentage</i> — Specifies the amber alarm threshold.
Values	1 to 1000

resv-cbs

Syntax	resv-cbs <i>percentage-of-pool-size</i> [amber-alarm-action step <i>percent</i> max <i>maximum</i>] no resv-cbs
Context	config>qos>named-pool-policy>q1-pools>pool
Description	The resv-cbs command is used to override the default reserved CBS size of the pool. The reserved CBS size defines the amount of buffer space within the pool that is not considered shared. When queues request buffers from the pool, they will be either within-CBS or above-CBS. If the queue is within-CBS based on the current queue depth and the configured CBS value for the queue, the requested buffer is taken from the reserved portion of the buffer pool. After the queues depth is beyond its configured CBS, the buffer will be taken from the pools shared space. Shared space buffers are subject to the WRED slope function within the buffer pool. If the WRED slopes are enabled, the buffer request may be denied based on WRED drop probability. The no form of the command restores the default reserved CBS size.
Default	resv-cbs 30
Parameters	<i>percentage-of-pool-size</i> — The <i>percentage-of-pool-size</i> parameter is required and is an integer specifying a percentage from 0 to 100 percent. Specifying a value of 30 (the default) is equivalent to specifying no resv-cbs. Values 0 to 100 Default 30 amber-alarm-action step <i>percent</i> — specifies the percentage step-size for the reserved CBS size of the pool. When using the default value, the adaptive CBS sizing is disabled. To enable adaptive CBS sizing, step <i>percent</i> must be set to a non-default value along with the max parameter. When reserved CBS is default, adaptive CBS sizing cannot be enabled. The reserved CBS (Committed Burst Size) defines the amount of buffer space within the pool that is not considered shared. Values 1 to 100 Default 0

max maximum — Specifies the maximum percentage for the reserved CBS size of the pool. When using the default value, the adaptive CBS sizing is disabled. To enable adaptive CBS sizing, **max** value must be set to a non-default value along with the **step percent**. When reserved CBS is default, adaptive CBS sizing cannot be enabled. The reserved CBS (Committed Burst Size) defines the amount of buffer space within the pool that is not considered shared. Max reserved CBS must not be more than the reserved CBS.

Values 1 to 100

Default 0

slope-policy

Syntax	slope-policy <i>slope-policy-name</i> no slope-policy
Context	config>qos>named-pool-policy>q1-pools>pool
Description	The slope-policy command is used to override the default slope-policy configuration for the named buffer pool. The specified <i>slope-policy-name</i> must exist as a current slope policy name. If the slope policy does not exist, the slope-policy command will fail. If a slope policy is currently associated with a named pool within a named pool policy, the slope policy cannot be removed from the system. The slope policy contains the High and Low WRED slope definitions that will be used by the pool on each XMA or MDA on which the pool is created. If the slope-policy command is not executed or the no slope-policy command is executed, the default slope policy will be associated with the pool. The no form of this command restores the default slope policy to the named pool.
Default	slope-policy "default"
Parameters	<i>slope-policy-name</i> — The slope-policy-name parameter is required and must specify an existing slope policy name. If slope-policy-name does not exist, the slope-policy command will fail.

port-allocation-weights

Syntax	port-allocation-weights default <i>weight mda weight port weight</i> no port-allocation-weights
Context	config>qos>named-pool-policy

Description The port-allocation-weights command is used to define the weights used to divide the buffers managed by a port into three categories: default, XMA and MDA, and port. The default category is given to the default pools, the XMA or MDA category is given to the XMA or MDA named pools, and the port category is used by the local port named pools. When the IOM is placed in named-pool-mode, each port has an inherent set of weights that places all port-managed buffers into the default category (default = 100, XMA or MDA = 0, port = 0). The policy port-allocation-weights command is used to override this port inherent behavior. When the policy is applied to the XMA or MDA, the defined port-allocation-weights parameter values override the inherent values for all ports on the XMA or MDA. When the policy is applied to the port level, the defined port-allocation-weights override both the local ports inherent weights and the XMA- or MDA-level named pool policy weights (if existing).

The **no** form of this command resets all values to the default value.

Parameters **default weight** — The default keyword is used to identify the weight value for the port where the policy is applied, and is used in the calculation of the amount of buffer space given to the default pools by the port. The following weight parameter is required and must be specified as an integer between 0 and 100. The specified weight only has meaning when compared to the mda and port weights. The sum of all three weights is divided into each weight to determine the amount of buffering given to the pools of each type.

Values 0 to 100

Default 50

mda weight — The mda keyword is used to identify the weight value for the port where the policy is applied, and is used in the calculation of the amount of buffer space given to the XMA- or MDA-level named pools by the port. The following weight parameter is required and must be specified as an integer between 0 and 100. The specified weight only has meaning when compared to the default and port weights. The sum of all three weights is divided into each weight to determine the amount of buffering given to the pools of each type.

Values 0 to 100

Default 50

port weight — The port keyword is used to identify the weight value for the port where the policy is applied, and is used in the calculation of the amount of buffer space given to the local port named pools by the port. The following weight parameter is required and must be specified as an integer between 0 and 100. The specified weight only has meaning when compared to the mda and port weights. The sum of all three weights is divided into each weight to determine the amount of buffering given to the pools of each type.

Values 0 to 100

Default 50

13.4.2.3 Show Commands

13.4.2.3.1 Named Pool Show Commands

named-pool-policy

Syntax	named-pool-policy named-pool-policy <i>policy-name</i> [detail] named-pool-policy <i>policy-name</i> association
Context	show>qos
Description	This command displays information on named pool policies.
Parameters	<i>policy-name</i> — Each named pool policy must be uniquely named within the system. Names of up to 32 ASCII characters composed of printable, 7-bit ASCII characters are supported. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. A named pool policy must exist prior to applying the policy name to an XMA, MDA, or port. detail — Displays detailed information on the given policy name. association — Displays the associations connected to the given policy name.
Output	The following output is an example of named pool policy information.

Sample Output

```
show qos named-pool-policy
=====
Named-Pool Policies
=====
Policy Name                Description
-----
test                        (not-specified)
test57                      (not-specified)
```

card

Syntax	card <i>card-slot</i> detail match "Named Pool Mode"
Context	show
Description	This command checks the card-specified named pool mode.
Output	The following output is an example of card named pool information.

Sample Output

```
show card 1 detail | match "Named Pool Mode"
=====
Named Pool Mode           : Configured (Enabled)
=====
```

mda

Syntax	mda slot[/mda] [detail]
Context	show
Description	This command displays named pool policies configured for an MDA.
Output	The following output is an example of named pool policies configured for an MDA information

Sample Output

```
show mda 1/2 detail
-----
QoS Settings
-----
Ing. Named Pool Policy      : test
Egr. Named Pool Policy      : test
=====
```

pools

Syntax	pools port-id
Context	show
Description	This command displays XMA, MDA, or port pools. If the pool size is zero, there are no queues associated with the pool and the pool is not in use; it is configured but not instantiated. To display details about an ingress or egress named pool, use the command show pools 1/2 [ingress egress] p2 . The output of the command shows which queues are using the named pool specified.
Output	The following output is an example of XMA, MDA, or port pool information.

Sample Output

```
*A:PE# show pools 1/2 ingress "p2"
=====
Pool Information
=====
MDA           : 1/2           FP           : 1
Application   : Ingress      Pool Name    : p2
```

```

CLI Config. Resv CBS : 30%           Policy Name           : 3pools
Resv CBS Step        : 0%           Resv CBS Max          : 0%
Amber Alarm Threshold: 0%           Red Alarm Threshold    : 0%
-----
Utilization           State          Start-Avg      Max-Avg      Max-Prob
-----
HiPlus-Slope         Down            85%           100%         80%
High-Slope           Down            70%           90%          80%
Low-Slope            Down            50%           75%          80%
Exceed-Slope         Down            30%           55%          80%
Time Avg Factor      : 7
Pool Total           : 149760 KB
Pool Shared          : 103680 KB      Pool Resv           : 46080 KB
-----

Current Resv CBS      Provisioned   Rising          Falling          Alarm
%age                  all Queues   Alarm Thd       Alarm Thd        Color
-----
30%                   0 KB        NA              NA              Green
Pool Total In Use    : 0 KB
Pool Shared In Use   : 0 KB
WA Shared In Use     : 0 KB
HiPlus-Slope Drop Pr*: 0           Hi-Slope Drop Prob : 0
Lo-Slope Drop Prob  : 0           Excd-Slope Drop Prob : 0
=====
Queue : 100->1/2/1->2
=====
FC Map                : af
Tap                   : 1/*          FP                : 1
Admin PIR             : 10000000     Oper PIR          : Max
Admin CIR             : 0           Oper CIR          : 0
Admin MBS             : 12240 KB     Oper MBS          : 12240 KB
High-Plus Drop Tail   : 12240 KB     High Drop Tail    : 12240 KB
Low Drop Tail         : 10800 KB     Exceed Drop Tail  : 9600 KB
CBS                   : 0 KB         Depth             : 0
Slope                 : not-applicable
=====
* indicates that the corresponding row element may have been truncated.
*A:PE#

```

port

Syntax	port <i>port-id</i> detail
Context	show
Description	This command displays named pool policies configured for a given port.
Output	The following output is an example of named pool policies configured for a given port information.

Sample Output

```
show port 1/2/10 detail
=====
Ing. Pool Policy   : n/a
Egr. Pool Policy   : test
=====
```

```
show port 1/2/10 detail
=====
Ing. Pool % Rate   : 100                               Egr. Pool % Rate : 100
=====
```

14 High Scale Ethernet MDA Capabilities

14.1 HSMDA QoS Model

This section describes QoS capabilities of the High Scale MDA (HSMDA). The HSMDA extends subscriber and service density of the first and second generation IOMs by adding an XMA or MDA level of ingress and egress queues, shapers and schedulers.

The HSMDA replaces the ingress and egress service queuing function performed by the ingress and egress forwarding plane on the IOM, providing up to 160K service or subscriber-based queues in each direction. The queues are configured in groups of eight. The addition of the ingress packet classification and queues allows QoS classification and service-level queuing to be performed on the XMA or MDA instead of the directly-attached IOM forwarding plane. The egress side of the HSMDA relies on the existing egress forwarding plane to map egress packets to the HSMDA egress queues.

The HSMDA moves service and subscriber-level ingress and egress QoS functions off the IOM forwarding plane to the XMA or MDA. The HSMDA QoS model provides the following features:

- Expanded queue scale for ingress and egress
- Hardware implemented provider style port-based scheduling
- Elimination of ingress dual-pass queuing at the ingress hardware
- Egress intermediate destination (such as DSLAM) shaping using secondary shapers
- Expanded counters at ingress and egress
- CIR bypass for color aware policing
- Ingress queue policing mode
- RED queue congestion control
- Egress dot1p remarking per packet, based on egress queue scheduling rate
- Per queue packet byte offset for queue stats, queue PIR, queue CIR, and queue group PIR accounting

14.1.1 Queue Scaling

The HSMDA supports 160,000 queues in the ingress and egress directions. Each queue supports two leaky buckets that perform a PIR shaping function and a CIR marking function based on scheduled rate out of the queue. Each queue also supports two RED slopes that may be used for managing queue congestion.


The queues are grouped into sets of eight queues each; 3840 queues are reserved for the system. A set of eight queues is called a queue group and is internally identified by a queue-group-id. Each queue within a group is numbered from 1 to 8, represented internally as the queue-id. The queue-id has an implicit scheduling class association based on the number of the queue (for example, queue 3 is a member of scheduling class 3). Queue groups are dynamically mapped to egress ports on an as-needed basis.

Individual queues are also mapped to 1 of 64 secondary shapers, each representing an egress intermediate destination (such as a DSLAM). One is pre-allocated per egress port, so only 54 are user-definable in 10x1G HSMDAs and 63 are user-definable in 1x10G HSMDAs.

14.1.2 Port-Based Scheduling


Forwarding for each egress port on the HSMDA is managed by a port-based scheduler. Each port-based scheduler maintains a maximum of eight strict forwarding levels (see [Figure 40](#)). Strict level 8 is the highest priority while strict level 1 is the lowest. There are also eight scheduling classes that contain each of the queues assigned to the port scheduler. A queue's membership in a scheduler class is controlled by the queue's identifier. For example, all queues with a queue-id equal to 1 are in scheduler class 1 while all queues with queue-id equal to 2 are in scheduler class 2. By default, each scheduler class is directly mapped to its corresponding strict scheduling level.

Figure 40 HSMDA Queue Mapping to Scheduler Class Service Lists



To allow for weighted servicing of selected scheduling classes, the port scheduler allows for two weighted groups to be optionally created (weighted-group-1 and weighted-group-2) and each may be populated with up to three consecutive scheduling classes (see [Figure 41](#)). The group itself maps to the highest inherent strict scheduling level of its member scheduling classes. Each scheduling class in a scheduling group are individually weighted, which allows for all queues represented by the class to be serviced according to the ratio of weights based on the active classes in the group (see [Figure 42](#)).

Figure 41 Scheduler Class Mapping to Strict Level or Weighted Group Example




OSSG149

Figure 42 Scheduler Weighted Group Configuration Example



OSSG150

Figure 43 Scheduler Class and Weighted Group Scheduling Priority Mapping Example



OSSG151

The port-based scheduler supports a port-based shaper used for creating a sub-rate condition on the egress port. Each individual strict scheduling level may also be configured with a shaping rate used to limit the amount of bandwidth allowed for that strict level, as shown in [Figure 43](#). In all, shaping PIRs can be defined at the following points in the queuing and scheduling architecture:

- Per port shaper
- Per strict level shaper
- Intermediate destination shaper (egress)

- Per queue-group shaper
- Per queue shaper

14.1.3 Dual Pass Queuing

In the standard queuing model, the ingress hardware performs both per service ingress SLA enforcement as well as per switch fabric destination-based virtual output queuing. Due to the requirement that each queue on ingress be mapped to a single switch fabric destination, when a SAP for a service type that forwards to multiple switch fabric destinations (such as VPLS, IES, and VPRN services), a single ingress service queue is created as multiple hardware virtual output queues in the hardware. In environments where the SAP or subscriber density requires more service queues than can be created with the available ingress hardware queues, two passes through the hardware are performed. This dual-pass mechanism allows a single hardware queue to represent a service queue on the first pass through the hardware while the second pass allows the hardware to use shared virtual output queues into the switch fabric.

HSMDA removes the need to perform two passes through ingress since the service or subscriber-based queuing is being performed on the XMA or MDA. This frees the ingress hardware to perform the virtual output queuing function using shared queues mapped to switch fabric destinations.

14.1.4 Egress Intermediate Destination Secondary Shapers

The HSMDA supports connection to an aggregation network servicing multiple DSLAMs or other subscriber last-mile aggregation devices such as DSLAMs or gigabit passive optical network optical network terminals (GPON ONTs). When operating in this fashion, multiple subscriber aggregators will be reached over an HSMDA 10GE port. The egress port scheduler combines all subscriber queues of the same scheduling class, and services the queues in a byte-fair-round-robin fashion. An effect of this behavior results in more packets being forwarded into the aggregation network towards a DSLAM than the DSLAM can accept. If the HSMDA egress port is congested, the egress bandwidth represented by the downstream discarded packets to the DSLAM may have been allocated to packets destined to

other DSLAMs. The HSMDA supports egress secondary shapers that are used to provide a control mechanism to prevent downstream overruns without affecting the class-based scheduling behavior on the port. When used in this manner, the secondary shaper is called an intermediate destination shaper. Egress secondary shapers are configured on a per port basis.


In other systems, the downstream destination is represented as a scheduler in a tiered hierarchical scheduler. This means that bandwidth is allocated on a per DSLAM basis without regard to the class of packets being forwarded to that DSLAM. In this model, a set of subscribers on one DSLAM receiving packets for a premium service are treated equally to subscribers on other DSLAMs receiving packets associated with best-effort type services.

The intermediate destination secondary shaper solves the downstream overrun issue without sacrificing class preference at the HSMDA egress port scheduler. All subscribers destined to the same DSLAM have their queue groups mapped to the same egress secondary shaper. As the scheduler services the queues within the groups according to their scheduler class, the intermediate destination shaper is updated. When the shapers rate threshold is exceeded, scheduling for all queues associated with the shaper is stopped. When the dynamic rate drops below the threshold, the queues are allowed to be placed back on to the scheduler service lists. By removing the queues from their scheduling context for a downstream congested DSLAM, the port scheduler is allowed to fill the egress port with packets destined to other DSLAMs without sacrificing class behavior on the port.

The system supports 64 egress secondary shapers that may be dynamically associated with an egress physical port. Each HSMDA queue must be associated with a secondary shaper. By default, a secondary shaper is created per physical port and all queues within a queue group on the port are associated with the default secondary shaper. All default secondary shapers will operate at the maximum rate and have no impact on queue scheduling. The hardware supports mapping individual queues within a queue group to separate secondary shapers. Based on the intended use case for secondary shapers, the provisioning model does not allow queues within the same queue group to map to separate secondary shapers. Queues are mapped to a secondary shaper on a queue group basis.

[Figure 44](#) shows HSMDA egress queue group and secondary destination shaper behavior.

Figure 44 HSMDA Egress Queue Group and Secondary Destination Shaper Behavior



OSSG152

14.1.5 Packet and Octet Counting

Each queue group supports a set of 16 counters. Each set of counters is identified by a counter-id and contains individual counters for:

- Discarded out-of-profile (low priority) packets

- Discarded out-of-profile (low priority) octets
- Discarded in-profile (high priority) packets
- Discarded in-profile (high priority) octets
- Forwarded out-of-profile packets
- Forwarded out-of-profile octets
- Forwarded in-profile packets
- Forwarded in-profile octets

The discard counters are incremented during enqueueing discarded events and the forward counters are incremented during scheduled dequeuing events.

For standard ingress queues, the offered stats-per-queue are counted by the ingress forwarding plane. For standard egress queues, the offered stats are derived by adding the hardware per-queue discard and forward stats. This means packets waiting to be scheduled in a standard egress queue will not be counted until it is forwarded. The HSMDA queue-offered stats operate the same as the standard egress queues where the discard and forwarded stats are combined to derive the offered stats for an HSMDA queue.

The decision on the counter to use is made per packet by HSMDA ingress hardware or by the egress forwarding plane hardware. The default behavior is to use the counter-id that corresponds to the queue-id to which the packet is mapped. This means the packets destined to queue 2 will be accounted for by counter set 2 within the same queue group. This sets aside the first eight counter sets as the default counters for the queue group. The remaining eight are available as counter override decisions.

A counter override can be performed within the ingress QoS classification rules wherever an HSMDA is installed.

The eight counter sets used as exception counters are identified as counter 1 through counter 8. While the discard and forwarding statistics can be overridden based on exception criteria, the offered statistics are maintained per queue. This means that the offered statistics for a queue includes all packets offered to the queue, but the discard and forwarding statistics only reflect packets handled by the queue that have not been associated with exception counters. It is possible to estimate the number of packets not represented by the queue statistics by subtracting the discard and forwarding statistics from the queue-offered statistics. The resulting number may be off slightly if packets are still in the queue when the statistics were collected, but this error is minimized when calculated over an appropriate amount of time and can be completely eliminated if the queue is allowed to drain prior to performing the calculation. When the queue statistics and the exception statistics are considered as a whole, all packets handled by the queue group are accurately represented by the counters.

Since the HSMDA only updates a single counter set per packet, overriding the counter for a packet causes that packet not to be represented within the default counter-id for the queue. If the queue counter-id is being used to determine CIR or PIR accuracy or basic throughput for a queue, any packets forwarded through the queue using a counter override are considered.

14.1.6 Above-CIR Discard with PIR Bypass

14.1.6.1 HSMDA Ingress Queue Policing Mode

Since HSMDA queue-based CIR and PIR leaky bucket behavior is driven by scheduling events from the queue, a true policing function where the ingress offered rate determines the color of the packet (green, yellow or red) is not available, by default. While enabling a PIR-based shaping rate for the queue will perform a similar function, the shaping function is just stopping the queue from scheduling. When the queue is stopped, packets are allowed to be buffered within the queue up to the RED slope or MBS-configured limits. The jitter for a shaping queue is based on how full the queue can get and how fast the queue is scheduled (influenced by the queue PIR, queue group PIR, and ingress secondary shapers). While the MBS for the queue may be configured to minimize jitter by preventing an excessive amount of data to accumulate in the queue, a policing mode can be enabled on the queue that uses the CIR leaky bucket to dynamically discard packet above the CIR threshold.

The HSMDA ingress queue policing mode behaves normally while the CIR leaky bucket is below its configured threshold. If the CIR fill depth rises above the threshold, the packet is discarded without updating any of the ingress schedulers PIR leaky buckets. The result is that while the queue is operating within its CIR, scheduled packets will be forwarded to the ingress forwarding plane for further processing and each packet will update all applicable PIRs. But when the queue scheduled rate rises above the CIR, scheduled packets are discarded. By discarding out-of-profile packets, the policing rate is enforced without unnecessary jitter based on queue congestion.

In order for ingress policing mode to function properly, the following configuration guidelines should be observed:

- The queue should be scheduled at the highest appropriate strict scheduling priority. If the queue is not scheduled at a high enough priority, scheduling from the queue may momentarily stall. The scope of this issue is limited since the ingress port bandwidth is less than the available scheduling bandwidth to the ingress forwarding plane.

- Ensure that the policing queue is not stalled by the queue groups configured aggregate rate limit. If the queue is not the highest scheduling priority, the sum of the allowed scheduling for the queues with higher scheduling priority may cause the queue groups PIR to be exceeded and thus scheduling for the policing queue will stall. If the queue is at a high enough priority, lower priority queues will only be allowed to consume the group PIR while the higher priority queues are inactive (empty). In the event that lower priority queues cause the PIR to suspend scheduling for the queue group, higher priority queues will have first access to ingress scheduling when the group PIR decrements below the threshold.
- Care should also be taken with ingress secondary shapers. If the queue is assigned to an ingress secondary shaper, the queue may be stalled when the aggregate rate of the queues associated with the shaper exceeds the shapers PIR.

Packets sent to an HSMDA ingress queue configured for policing cannot have the ignore-CIR flag set. Color aware profiling and ingress queue policing should not be mixed. In the event that a packet is classified as in-profile or out-of-profile while ingress queue policing is configured, the ignore-CIR bit will automatically be reset to zero (the queue CIR will be updated by the packet).

Two implementation options exist to account for out-of-profile scheduling discards:

1. When a packet is discarded due to out-of-profile scheduling, the out-of-profile packet and octet forward counters within the counter ID associated with the packet are updated. If the packet is forwarded, the in-profile packet and octet forward counters within counter ID' are incremented. Software must add the out-of-profile forwarded counter to the low-priority discard counter to determine the total discards based on out-of-profile for the queue.
2. Alternatively, the discard event may be hard-coded to increment the discard counters directly. When discarding, the congestion-priority bit is used to determine whether the high or low discard counter is incremented.

For packets discarded at enqueueing time, the high-priority or low-priority packet and octet discard counters are updated. The decision to use the high- or low-priority discard counter is driven by the congestion-priority' bit associated with the packet. This bit is set based on both color aware profiling and ingress priority of the packet. If the packet is explicitly classified as in-profile, the bit is set to high. If the packet is explicitly classified as out-of-profile, the bit is set to low. If the profile of the packet is undetermined (not explicitly in-profile or out-of-profile), the bit is set to high or low based on the classified ingress priority of the packet.

14.1.7 HSMDA Buffer Utilization Controls

The HSMDA has 1 million ingress and 1 million egress 168-byte buffers available for packet queuing purposes; an average of approximately six buffers per queue when all 163,480 are active. Certain queues need more than 6 buffers while other queues require very shallow buffering based on the type of traffic the queue is servicing and the scheduling priority of the queue. To facilitate management of the available buffer space, the HSMDA supports a hierarchical buffer pool scheme and a per queue set of RED slopes. The buffer pools allow proper sharing of the buffer space while the slopes within each queue set limits on how many buffers each queue may consume.

14.1.7.1 HSMDA Buffer Pools

Two types of queues are created on the HSMDA: provisioned service or subscriber queues and system-created queues. System queues are transparent to the user and perform functions like discard bypass. Since system queues are critical to the operation of the system, normal service or subscriber queue activity will not cause buffer starvation on the system queues. Buffer utilization is separated based on the scheduling class, ensuring that activity on one set of class queues does not impact buffer availability for other class queues.

The pooled buffer management capabilities in the HSMDA include:

- Identifying which queue groups consume provisioning buffers and which consume system reserved buffers.
- Setting the total provisioning buffers available per port for each scheduling class.
- Setting the total system reserved buffers available per port for each scheduling class.
- 32 aggregation buffer pools used for managing buffers available per class and per type (provisioned and system).

14.1.7.2 Identifying Queue Groups as Provisioned or System

All queues are contained in a set of eight queues called a queue group. Two sets of 20,480 queue groups exist on the HSMDA, one set for ingress and the other for egress. The queue groups are defined as either provisioned (service or subscriber) or reserved for system use. The HSMDA uses a two 20,480 bit-wide tables to allow the system to define each queue group as either provisioned or system reserved separately for ingress and egress. The queue group id mapping table represented in

Figure 45 shows queue groups in pool group P (provisioned) or pool group S (system). For ingress and egress, the first 20,000 table entries (0 to 19,999) are set to group P and the remaining 480 (20,000 to 20,479) are set to group S.

Figure 45 Queue Group ID Mapping Table

0	1	2	3	4	5	6	7	8	9	10	11				20479
P/S	P/S	P/S	P/S	P/S	P/S	P/S	P/S	P/S	P/S	P/S	P/S	•	•	•	P/S

OSSG154

14.1.7.3 Provisioned and System Port Class Pools

The HSMDA uses a second table called the port class buffer pools table (see Figure 46) that represents a set of 160 buffer pools. Pools 0 through 79 are used by queues within group P and 80 through 159 are used by queues within group S. Each set of eighty pools is divided into 10 subsets of 8 pools each. Each subset is dedicated to a physical port on the HSMDA. For group P, pools 0 through 7 are for port 1, 8 through 15 are for port 2, and so on, up to pools 72 through 79 for port 10. A queue is mapped to a pool based on the queue-group-id mapping to P or S, then the port the queue is associated with is used to pick the subset. Within the subset, the internal queue-id is used as an offset to pick an actual pool. Queue-id 0 (provisioned as 1) on port 3 in group P is mapped to pool 16. Each pool in the table also has two aggregate pool pointers used to provide further control on buffer allocation. Agg-Pool-Ptr-1 and Agg-Pool-Ptr-2 arbitrarily tie the port class pool to two aggregate pools from a third table of 32 buffer pools.

Figure 46 Port Class Buffer Pools Table

0	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	Group P
1	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	
2	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	
3	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	
⋮				
79	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	Group S
80	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	
81	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	
82	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	
83	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	
⋮				
159	Buffers-Available (21bits)	Agg-Pool-Ptr-1 (5bits)	Agg-Pool-Ptr-2 (5bits)	

OSSG155

14.1.7.4 Aggregate Pools for Type and Class Separation

The third table is called the aggregate control buffer pools table (see [Figure 47](#)). The table consists of 32 buffer pools that may be used arbitrarily by the port class pools. While the association from port class pool to aggregate control pool is arbitrary based on the 5-bit pointers, it is expected that the control pools will be divided into two groups of 16 pools, each group having two sub-groups of eight pools each. The first aggregate control group (pools 0 through 15) will be for provisioned buffer management and will be used by group P port class pools. The second aggregate control group (pools 16 through 31) will be for system-level buffer management and used by group S port class pools.

Figure 47 Aggregate Control Buffer Pools Table

0	Buffers-Available (21bits)	}	Group P
1	Buffers-Available (21bits)		
2	Buffers-Available (21bits)		
3	Buffers-Available (21bits)		
4	Buffers-Available (21bits)		
⋮			
15	Buffers-Available (21bits)	}	Group S
16	Buffers-Available (21bits)		
17	Buffers-Available (21bits)		
18	Buffers-Available (21bits)		
19	Buffers-Available (21bits)		
⋮			
31	Buffers-Available (21bits)		

OSSG156


14.1.7.5 Use of Aggregate Control Buffer Pools

The aggregate control buffer pools are separated into two sets. The first 16 pools (0 through 15) are used by the provisioned group (group P) port class pools. The second 16 pools (16 through 31) are used by the system group (group S) port class pools.

The first 8 (0 through 7) are used as class based pools. Pool 0 is used by scheduling class 1 (0 internally), and up to pool 7 used by scheduling class 8. Agg-Pool-Ptr-1 for port class pools 0, 8, 16, 24, 32, 40, 48, 56, 64, and 72 (all port call pools associated with queues with queue-id 1) is set to aggregate control pool 0. This ensures that all provisioned queues in scheduling class 1 are limited based on the amount of buffering for class 1 on their port and also the total buffers used by the class is limited for the XMA or MDA. If either of the pools is exhausted, the queue will not receive a buffer. In like manner, Agg-Pool-Ptr-1 for all port class pools 1, 9, 17, 25, 33, 41, 49, 57, 65, and 73 is set to pool 1. This is true up to port class pools 7, 15, 23, 31, 39, 47, 55, 63, 71 and 79 having Agg-Pool-Ptr-1 set to pool 7.

The second 8 aggregate control buffer pools (8 through 15) are used as provisioned root pools. The purpose of the root pools is to allow the class pools to be oversubscribed without the possibility of the provisioned buffer usage stealing buffers from the system reserved buffers. Before sizing the provisioned root pools, a portion of the total buffer space is set aside for system purposes. The remaining buffers are divided between the provisioned root pools based on a weight parameter in each root pool. The weights may be set between 0 and 100. A value of zero indicates that a specific provisioned root pool will not receive buffers (pool size will be 0). Because root pools cannot be oversubscribed, they provide a protection mechanism for higher level pools. The number of root pools in use is dependent on the HSMDA pool policy applied to the XMA or MDA (ingress and egress are controlled by independent policies). The aggregate control class pools are associated with the root pools through the policy as well. [Figure 48](#) shows the buffer pool hierarchy.

Figure 48 Buffer Pool Hierarchy



OSSG157

14.1.7.6 HSMDB Buffer Pool Policy

HSMDB buffer pool policies contain information the system uses to configure the individual pool sizes and the root pools used by the class pools. The division between total provisioned and total system buffers is based on the system-reserve parameter that defines the percentage of buffers reserved for system use. The remaining buffers are placed in the provisioned root buffer pools based on the weight associated with each root pool.

The eight class pools within the policy are defined with a root-pool parent association (1 through 8) and a percentage value. The parent associated with the class-pool defines the Agg-Pool-Ptr-2 setting in the port class pools table for each port class pool using the class-pool. For example, if class-pool 1 (aggregate control pool 0) is associated with root-pool 3 (aggregate control pool 10) within the policy, port class pools 0, 8, 16, 24, 32, 40, 48, 56, 64, and 72 will contain a value of 10 in Agg-Pool-Ptr-2. The percentage used to size the class pool is applied to the class pool's root-pool size to derive the class pool size. The sum of the sizes of the class pools parented by a root-pool may exceed (oversubscribe) the root-pool size.

System pools are managed similarly. The actual system port class pools and system aggregate control pools (16 through 31) are not user configurable.

14.1.7.6.1 Default HSMDB Buffer Pool Policy

An HSMDB buffer pool policy named **default** always exists on the system and cannot be deleted or edited. The default policy is used for ingress and egress on all HSMDBs until an explicitly created HSMDB policy is defined on the HSMDB.

The default policy parameters are listed in [Table 111](#).

Table 111 Default Policy Parameters

System Reserve	
Percentage:	10%
Root Pools:	
Root-Pool 1 Weight:	75
Root-Pool 2 Weight:	25
Root-Pool 3 Weight:	0
Root-Pool 4 Weight:	0
Root-Pool 5 Weight:	0

Table 111 Default Policy Parameters (Continued)

System Reserve	
Root-Pool 6 Weight:	0
Root-Pool 7 Weight:	0
Root-Pool 8 Weight:	0

The class pool parameters are listed in [Table 112](#).

Table 112 Class Pool Parameters

Class Pools		
Class-Pool 1	Parent:	Root-Pool 1
	Percentage:	40%
Class-Pool 2	Parent:	Root-Pool 1
	Percentage:	35%
Class-Pool 3	Parent:	Root-Pool 1
	Percentage:	30%
Class-Pool 4	Parent:	Root-Pool 1
	Percentage:	25%
Class-Pool 5	Parent:	Root-Pool 1
	Percentage:	20%
Class-Pool 6	Parent:	Root-Pool 2
	Percentage:	50%
Class-Pool 7	Parent:	Root-Pool 2
	Percentage:	40%
Class-Pool 8	Parent:	Root-Pool 2
	Percentage:	30%

14.1.7.7 Port Class Pool Sizing

The port class pools are sized based on the port's active bandwidth, provisioned use of the port, and the port bandwidth rate modifier percentages defined on the port. For the HSMDB, a port's active bandwidth is the current speed provisioned for the port. For ingress, this is the current line rate of the port. For egress, it is the lesser of the port's line rate and the port scheduler's current maximum rate. The active bandwidth of the port is considered to be zero when a SAP or subscriber has not been provisioned on the port. When a queue group is associated with the port, the actual active bandwidth of the port is used for port class buffer pool sizing.

The active bandwidth can be modified by the **max-rate** commands on the port. The parameters are used to artificially increase or decrease the amount of buffers that may be used by the port. These commands have no effect on the actual bandwidth used by the port.

The system uses the active bandwidth of each port to decide on how much each class pool the port should receive. The port's active bandwidth is divided into the sum of all ports active bandwidth to derive the port's pool factor. This factor is multiplied by the size of each class pool and the result is applied to the port's class pool. Thus, the sum of the sizes of a given port class pool over all ports should equal the size of the actual aggregate control class pool.

14.1.7.8 HSMDB Available Buffer Register Operation

The HSMDB considers a buffer pool empty when the buffer-available register for the pool drops below 64. As buffers are allocated to a queue, the HSMDB evaluates whether the queue's port class pool buffer-available register is less than 64. If not, the HSMDB then looks at the aggregate control pools pointed to by the port class pool's Agg-Pool-Ptr-1 (class-pool) and Agg-Pool-Ptr-2 (root-pool) to determine whether either of the pool's buffer-available registers are less than 64. If not, then the buffers required to enqueue the packet on the queue are given to the queue and the buffer-available registers are decremented based on the number of buffers given. As packets are scheduled out of the queue, the buffers are returned to the free list and the buffer-available registers are incremented by the correct amount.

The HSMDB stops allocating buffers at 63 or fewer remaining to allow packet size fairness on the pools. If it just stopped at 0, large packets would be at a disadvantage as the buffer pool neared the zero mark. If the register read 20 buffers remaining and a packet arrives needing 21 buffers, the pool would deny the buffer request (it would decrement below 0). But if a smaller packet arrives needing fewer buffers than 20, it would be allowed. And the buffer pool would continue to allow the smaller packets until the pool was depleted. By stopping at 63 or fewer, every packet has access to

the remaining pool since the maximum size packet requires fewer than 64 buffers. While this scheme allows for fair access to the buffer pool, some buffers will not be used by the HSMDA. This is considered inconsequential due to the limited number of buffer pools. The absolute worst case in buffer inefficiency is 12,096 buffers (63 buffers * 192 pools) out of 1 million buffers (about 1.2%).

When the pools are first initialized, each pool's buffer-available register is set to the number of buffers available on the pool. When the HSMDA buffer pool policy managing a pool is changed so that the number of buffers managed by the pool increases or decreases, the system will increment or decrement the current value of the buffer-available register based on the change in the pool size. The register supports a negative value for the case where the buffer pool size is decreased and the buffer-available register is currently less than the size of the decrease. If the buffer pool goes negative (or below 64 available buffers), no buffers are allocated by the pool until buffers associated with the pool are returned to the free list and the register increments to a value of 64 or higher.

14.1.8 HSMDA Queue Congestion and Buffer Utilization Controls

Each queue supports a 10 bit index into an HSMDA slope policy table. Each policy in the table consists of two RED slopes (high priority and low priority) for the purpose of managing queue congestion. Due to the large number of queues supported on the XMA or MDA, the ability of a queue to effectively manage a weighted sliding window of queue utilization is not practical. HSMDA RED slopes operate as on the instantaneous depth of the queue.

Each slope policy within the HSMDA consists of two slope definitions, each represented by a 14 bit start slope value, a 14 bit end slope value, and an 8-bit fixed point inverse slope value. The fixed point value is represented as a 4-bit whole value (values from 0 to 15) and a 4-bit fraction (from 0 to 0.9375 in 0.0625 increments).

In operation, a packet attempting to enter a queue triggers a check to see if the packet should be allowed based on queue congestion conditions. The packet contains a congestion-priority flag in the shim header telling the HSMDA whether to use the high or low slope. The slope policy containing the slope is derived from the policy index in the queue configuration parameters on the HSMDA. The XMA or MDA retrieves the current queue depth in buffers and the slope's three configuration values (start buffer, end buffer, and inverse slope value). Logically, the following algorithm is used to determine whether the packet should be allowed in the queue based on discard probability:

- If the queue depth is greater than or equal to the slope's end-buffer value, the packet is discarded.
- A random number in the range of 0 to 127 is generated.
- The random number is multiplied by the inverse slope value, then added to the slope Start-Buffer to derive the random fill depth that is the number of buffers that need to be full in order for the slope discard threshold to cross the random number.
- If the slope fill depth is equal to or greater than the random fill depth, the packet is discarded.
- The packet still may be discarded if a buffer pool associated with the queue has decremented to the discard point or if the free buffer list is exhausted.


14.1.8.1 Maximum HSMDB Queue Depth

Since RED slope discards are done based on the current queue depth before allowing a packet into the queue, a queue may consume buffers beyond the configured MBS value based on the size of the packet. The maximum number of buffers is based on the high slopes end-buffer parameter plus the maximum packet size less one buffer. When the slopes end-buffer is reached or exceeded, all other packets reaching the queue and associated with the slope will be discarded. When scheduling removes packets from the queue, the queue depth will decrease, eventually lowering the depth below the end-buffer threshold.

14.1.8.2 Control Plane HSMDB RED Slope Policy Management

RED slope configuration is managed by defining up to 1,024 named HSMDB slope policies on the chassis and mapping the queue to a specific policy name. Each slope policy contains configuration information for the high-priority slope and the low-priority slope. HSMDB Slope policies differ from standard slope policies in that they do not support the time-average-factor feature used to manage the weighting utilization of the buffer space the standard slopes are managing. HSMDB queue slopes operate based on instantaneous queue utilization and do not maintain a weighted utilization value. [Figure 49](#) shows the high- and low-priority RED slopes used to derive the discard probability based on the current depth of the queue.

Figure 49 High and Low RED Slopes



14.1.8.3 HSMDB Slope Policy MBS Parameter

Unlike standard queues, HSMDB queues do not have a configured maximum buffer size (MBS) parameter. Instead, the high and low RED slopes are used for all queue congestion control functions by the queue. The system uses an MBS value to define how much buffering an HSMDB queue may use, but it is contained in the HSMDB slope policy. The slope policy uses the configured MBS value to provide context for the slope parameters that are defined as percentages of MBS.

14.1.8.4 HSMDB Slope Policy Slope Parameters

Each slope in the slope policy is defined by three points: the starting-depth, slope maximum-depth, and the slope maximum-discard-probability. Only three points are required since the starting-depth value implies a starting-discard-probability of zero percent. The starting-depth defines at which queue depth the slope starts to rise from zero. The maximum-depth point defines the queue depth where the slope ends and goes straight to 100%. The maximum-discard-probability defines the how high the slope rises from zero before going straight to 100% at the maximum-depth point.

If the starting-depth and maximum-depth percentages are equal, the system performs a drop tail; the discard probability slope is essentially non-existent in this case.

The system takes the configured slope parameters and uses them to calculate the HSMDB internal slope definitions:

- The MBS value is multiplied by the starting-depth percentage to derive the number of bytes at which the slope will start. This value is then converted to buffers by dividing the resulting byte value by the buffer size (168 bytes) and rounding to the nearest integer value. This is the HSMDA slope start-buffer parameter (14 bit value).
- The same process is performed for the maximum-depth percentage to derive the HSMDA slope end-buffer parameter (14 bit value).
- The system calculates the Inverse-slope value using the following formula:
 - $\text{Inverse-slope} = (\text{end-buffer} - \text{start-buffer}) / \text{maximum-discard-probability} \times 1.27.$
- The Inverse-Slope value is converted to an 8 bit fixed point notation where:
 - Bits 7 through 4 represent the value above zero and bits 3 through 0 represent the fraction below zero.
- If the inverse slope is less than 0.0625, the system uses a value of 0.0625 (0000.0001) as the 8 bit value.
- If the inverse slope is greater than 15.9375, the system uses 15.9375 (1111.1111) as the 8 bit value.
- [Table 113](#) can be used to derive the 8-bit value for inverse slopes within the range 0.0625 and 15.9375.

Table 113 HSMDA Inverse Slope Fixed Point Binary Values

Inverse Slope Fixed Point Binary Format ($2^3 \ 2^2 \ 2^1 \ 2^0. \ 2^{-1} \ 2^{-2} \ 2^{-3} \ 2^{-4}$)	
Whole Number Decimal Value	Corresponding Binary Value ($2^3 \ 2^2 \ 2^1 \ 2^0$)
0	0000
1	0001
2	0010
3	0011
4	0100
5	0101
6	0110
7	0111
8	1000
9	1001

Table 113 HSMDA Inverse Slope Fixed Point Binary Values (Continued)

Inverse Slope Fixed Point Binary Format ($2^3 \ 2^2 \ 2^1 \ 2^0 \ . \ 2^{-1} \ 2^{-2} \ 2^{-3} \ 2^{-4}$)	
10	1010
11	1011
12	1100
13	1101
14	1110
15	1111
Fractional Number Decimal Value	Corresponding Binary Value ($2^{-1} \ 2^{-2} \ 2^{-3} \ 2^{-4}$)
0.0000	0000
0.0625	0001
0.1250	0010
0.1875	0011
0.2500	0100
0.3125	0101
0.3750	0110
0.4375	0111
0.5000	1000
0.5625	1001
0.6250	1010
0.6875	1011
0.7500	1100
0.8125	1101
0.8750	1110
0.9375	1111

For instance, if the MBS value is defined as 16,800 bytes and the low slope was configured with a starting depth set to 75%, a maximum depth set to 100% and a maximum discard probability set to 80%:

- The system takes 75 percent of 16,800 bytes and derives a starting slope at a queue depth of 12,600 bytes.
- The system takes 100 percent of 16,800 bytes and derives an ending slope of 16,800 bytes.
- The system converts the starting, ending, and slope duration values to number of buffers, then calculates the slope duration (end – start).
 - $12,600 / 168 = 75$ start-buffer
 - $16,800 / 168 = 100$ end-buffer
- The actual step of the slope may be calculated by dividing the maximum discard probability value with the slope run:
 - $80 / (100 - 75) = \text{slope step is } 3.2$ (for every buffer beyond the start of slope, the probability rises 3.2%, after 25 buffers the slope reaches 80%).
- But the HSMDA uses the inverse slope within its algorithm and a drop probability random number between 0 and 128. The system calculates the inverse slope by dividing the slope run by the maximum discard probability multiplied by 1.27 (converting 0 to 100 to 0 to 127):
 - $(100 - 75) / 80 * 1.27 = 0.396875$ inverse-slope
- The system takes the inverse slope step and converts it to an internal HSMDA fixed point binary notation with the most significant 4 bits representing the whole portion of the inverse slope (above 0) and least significant 4 bits representing the fractional portion of the slope (below 0).
 - The inverse slope is less than 0 so the most significant 4 bits are 0000.
 - The fractional portion of the inverse slope is 0.396875 and the closest result in four bits based on [Table 113](#) is a least significant bit value of 0110 (0.375 decimal).
 - The system concatenates the two results into an 8 bit number resulting in 0000110 as the inverse slope binary value.

The system uses the starting buffer value, the ending buffer value, and the inverse slope eight bit value to populate a slope definition into the HSMDA. Two slopes are populated per slope policy. Each slope policy is given an HSMDA slope index between 0 and 1023. Since every packet received on an HSMDA queue is associated with either the high or low slope, the provisioned MBS value is not required and is not a managed parameter for HSMDA queues.

14.1.8.5 HSMDB Slope Shutdown Behavior

HSMDB slope policies allow the high or low RED slope to be shutdown. This effectively configures the HSMDB internal slope to have a start buffer and end buffer equal to the buffer value closest to the configured policy MBS value. Packets matching the slope will still be discarded based on the slope, but it will appear that the MBS value is being enforced within the queue. The **maximum-discard-probability** is set to 100% when the slope is shutdown. The **inverse-slope** value sent to the XMA or MDA is set to all zeros (00000000) but has no effect since the starting and ending buffers are the same.

14.1.8.6 Ingress Packet Mapping to HSMDB RED Slope

At ingress, the following is used to determine packet mapping to a RED slope:

High slope

- Explicit profile classified as in-profile
- Undefined profile, but classified as high priority

Low slope

- Explicit profile classified as out-of-profile
- Undefined profile, but classified as low priority

14.1.8.7 Egress Packet Mapping to HSMDB RED Slope

At egress, the following is used to determine packet mapping to a RED slope:

High slope

- Egress reclassified highplus or high priority
- Non-reclassified in-profile

Low slope

- Egress reclassified low or exceed priority
- Non-reclassified out-of-profile

A default HSMDB slope policy for HSMDB queues with default parameters for the low and high slopes will be used to mimic the non-HSMDB queue drop tail behavior:

MBS	16,800 bytes (100 buffers)
High Slope	Start-depth 100% Max-depth 100% Max-probability 100% Shutdown
Low Slope	Start-depth 90% Max-depth 90% Max-probability 100% No shutdown

14.1.8.8 HSM DA Queue Congestion or Pool Congestion Discard Stats

When a packet is discarded, the high-priority or low-priority packet and octet discard counter is incremented. Which counter is incremented is based on the counter ID and the value of the congestion-priority flag within the packets HSM DA header. The classification rules within the SAP ingress QoS policy use the following logic to determine the discard counter associated with a packet:

Explicit in-profile (color aware profiling)	high priority
Explicit out-of-profile (color aware profiling)	low priority
Non-profiled, high priority	high priority
Non-profiled, low priority	low priority

The SAP egress QoS policy uses the following logic to determine the discard counter associated with a packet:

In-profile at ingress or egress reclassified highplus or high priority	high priority
Out-of-profile at ingress or egress reclassified low or exceed priority	low priority

The discard counters and the RED slope determination are both driven by the same classification results. When a packet is defined as explicitly inplus/in-profile or exceed/out-of-profile, the high or low priority of a packet is ignored at ingress.

14.1.9 Egress Queue CIR-Based Dot1p Remarking

The HSMDA adds the capability to perform remarking of one dot1p value within a dot1q or qinq-tagged packet, based on the dynamic CIR state of the egress queue at the time the packet is scheduled out the egress port. This allows downstream aggregation Layer 2 aggregation devices to manage congestion based on the dot1p field (including the DEI bit).

14.2 SAP Ingress and SAP Egress QoS Policies

The queue definition and scheduling behavior for HSMDA queues require different provisioning behavior from the standard QChip-based service-level queuing.

14.2.1 SAP Ingress QoS Policy

The SAP Ingress QoS policy performs three distinct functions:

- Service queue definitions
- FC and subclass queue mappings and ingress attributes
- Definition of hierarchical packet classification rules

For standard queuing, the application of the SAP ingress QoS policy to a SAP results in a set of hardware queues being dynamically assigned to the SAP representing the service queues defined in the policy.

The remaining sections of the policy affect resources within the hardware forwarding plane. Each ingress policy within the ingress forwarding plane consists of a set of classification rules consisting of dot1p, IP precedence, and IP DSCP tables used to match packets that ingress the SAP. Each table entry maps the packet to a forwarding class (each subclass is contained within a forwarding class, so subclasses indirectly map to a forwarding class). The section of the policy that represents the forwarding class to queue mappings is also represented by a table that allows the forwarding plane to take the forwarding class and determine which service queue (and the resulting hardware queue) will handle the packet.

The SAP ingress policy has been expanded so a single ingress policy contains queue definitions for both standard service queues and for HSMDA service queues. This provides a policy assignment model that does not need to know the difference between a SAP using standard service queues and another SAP using HSMDA service queues.

The standard ingress service queues are separated into two types: point-to-point and multipoint. Point-to-point queues are used either for VLL services such as Epipe or for unicast traffic within a VPLS, VPRN, or IES service. Multipoint queues are used by VPLS, VPRN, and IES services for packets that must be replicated by the switch fabric. Within the IES and VPRN services, only IP multicast routed packets are forwarded through the multipoint queues. Within VPLS services, broadcast (MAC DA = ff:ff:ff:ff:ff:ff), multicast (non-broadcast destination with multicast bit set), and packets with an unknown destination MAC address are forwarded through the multipoint queues. The standard service queues within the SAP ingress QoS policy

are numbered 1 through 32, of which 8 may be point-to-point queues and 24 multipoint queues. Since the HSMDA queues are not required to handle multipoint forwarding into the switch fabric, a distinction between point-to-point and multipoint is not present for HSMDA queues. Also, HSMDA queues are not created or dynamically assigned to a SAP or subscriber context. Instead, eight queues (numbered 1 through 8) always exist for each queue group on the XMA or MDA. Since the queue group is assigned to the ingress SAP, all eight queues within the group automatically are available for the SAP. This means that although a SAP ingress QoS policy does not reference a particular queue ID, that queue is available for forwarding class mappings.

Another difference between standard service queues and HSMDA service queues is that they cannot be associated with a port- or service-level virtual scheduler and so the port-parent and parent commands are not available. Instead, each queue is implicitly mapped directly to the HSMDA port or ingress scheduler based on the queue ID. Queue 1 is mapped to scheduling class 1 and queue 8 is mapped to scheduling class 8. Scheduling class 8 has the highest priority and class 1 has the lowest, unless a scheduling class is grouped with another scheduling class, in which case the group inherits the scheduling priority of the highest class within the group (the classes within the group are handled based on the weight assigned to each class).

The last major difference between standard service queues and HSMDA service queues is the support for RED slopes within the HSMDA queues. The HSMDA uses the slope policy that the queue is associated with to configure the contour of the high and low slope within the queue.

The SAP ingress QoS policy classification rule actions are also modified to allow for a counter override capability. The counter override function allows for one of the eight extra sets of counters per queue group to be used instead of the per queue counters. This is intended for diagnostic or exception-based accounting purposes.

14.2.2 SAP Egress QoS Policy

The SAP egress QoS policy requires the same type of modifications as the SAP ingress policy. The policy supports queue definitions for both standard service queues and HSMDA service queues and the forwarding class mappings to the individual queue IDs.

Another modification is the ability to define egress HSMDA counter override criteria that relies on an egress TCAM lookup based on IP flow criteria match entries. Egress IP flow-based HSMDA counter overrides are ignored when applied to a SAP not on an HSMDA. Egress counter overrides are ignored when a SAP is a member of an efficient multicast group.

14.3 Subscriber Queuing Differences

Standard service queues are instantiated on a SAP or subscriber sla-profile basis. Each SAP or subscriber sla-profile instance within a SAP has its own set of queues that are managed into an aggregate SLA by a software-based virtual scheduler. Due to its internal architecture, subscriber queues on the HSMDA are handled differently. The subscriber aggregate rate limit is represented by the queue group shaper. Since the queue group shaper only manages the eight queues within the group, all subscriber hosts within a subscriber context must share the same eight queues. Each subscriber SLA-profile (the object that allows different groups hosts for a single subscriber) must classify packets to the subscriber level queues and not to queues specific to the SLA-profile instance. While each SLA-profile instance does not maintain its own set of queues on the HSMDA, the packet classification rules are still maintained per sla-profile instance and not restricted to the subscriber level.

To allow provisioning of the sla-profile QoS classification rules and also provide the ability to specify the queuing behavior at the subscriber profile level, the SAP ingress and SAP egress QoS policies are reused in each case. At the sla-profile level, the SAP QoS policies packet forwarding class classification rules and forwarding class to queue ID mappings are in effect (the queue definitions are ignored on the sla-profile instances for HSMDA subscribers). At the subscriber profile level, the queue-id definitions are in effect while the packet classification rules and mappings are ignored.

14.4 Basic HSMDA Configurations

14.4.1 HSMDA Pool Policies

The following displays details of the **default** HSMDA pool policy configuration.

```
A:ALA-48>config>qos>hsmda-pool-policy# info detail
-----
no description
system-reserve 10.00
root-tier
  root-pool 1 allocation-weight 75
  root-pool 2 allocation-weight 25
  root-pool 3 allocation-weight 0
  root-pool 4 allocation-weight 0
  root-pool 5 allocation-weight 0
  root-pool 6 allocation-weight 0
  root-pool 7 allocation-weight 0
  root-pool 8 allocation-weight 0
exit
class-tier
  class-pool 1 root-parent 1 allocation-percent 40.00
  class-pool 2 root-parent 1 allocation-percent 35.00
  class-pool 3 root-parent 1 allocation-percent 30.00
  class-pool 4 root-parent 1 allocation-percent 25.00
  class-pool 5 root-parent 1 allocation-percent 20.00
  class-pool 6 root-parent 2 allocation-percent 50.00
  class-pool 7 root-parent 2 allocation-percent 40.00
  class-pool 8 root-parent 2 allocation-percent 30.00
exit
-----
A:ALA-48>config>qos>hsmda-pool-policy#
```

14.4.2 HSMDA Scheduler Policies

HSMDA scheduler policies can be assigned to an egress HSMDA port or as the ingress control scheduler between the HSMDA and the ingress forwarding plane. The policy contains the needed commands to provision the scheduling behavior of a set of HSMDA scheduler classes. When assigned to an HSMDA egress port, the policy is used to define the scheduling behavior for all queues associated with the egress port. When assigned to the ingress path of an HSMDA, the policy is used to define the scheduling behavior for all ingress queues on the HSMDA (regardless of ingress port).

The following displays details of the **default** HSMDA scheduler policy configuration:

```
*A:ALA-48>config>qos# info detail
```

```
#-----
echo "HSMDB Scheduler Policies Configuration"
#-----
    hsmdb-scheduler-policy "default" create
        description "Default hsmdb scheduler QoS policy"
        no max-rate
        group 1 rate max
        group 2 rate max
        scheduling-class 1 rate max
        scheduling-class 2 rate max
        scheduling-class 3 rate max
        scheduling-class 4 rate max
        scheduling-class 5 rate max
        scheduling-class 6 rate max
        scheduling-class 7 rate max
        scheduling-class 8 rate max
    exit
    hsmdb-scheduler-policy "HSMDB-test" create
        description "HSMDB policy"
        no max-rate
        group 1 rate max
        group 2 rate 50000
        scheduling-class 1 rate max
        scheduling-class 2 group 2 weight 1
        scheduling-class 3 rate max
        scheduling-class 4 rate max
        scheduling-class 5 rate max
        scheduling-class 6 rate max
        scheduling-class 7 rate max
        scheduling-class 8 rate max
    exit
#-----
...
*A:ALA-48>config>qos#
```

14.4.3 HSMDB Slope Policies

The following displays details of the **default** HSMDB slope policy configuration.

```
A:ALA-48>config>qos#
#-----
...
    hsmdb-slope-policy "default" create
        description "Default hsmdb slope policy."
        queue-mbs 16800
        high-slope
            start-depth 100.00
            max-depth 100.00
            max-prob 100.00
            no shutdown
        exit
        low-slope
            start-depth 90.00
            max-depth 90.00
            max-prob 100.00
```

```

        no shutdown
    exit
exit
...
-----
A:ALA-48>config>qos#

```

14.4.4 Egress Queue Group

The following displays details of the **default** HSM DA egress queue group configuration.

```

config
  qos
    queue-group-templates
      egress
        queue-group group-name
          hsm da-queues
            low-burst-max-class class
            packet-byte-offset {add add-bytes | subtract sub-bytes}
            wrr-policy wrr-policy-name
            queue queue-id [create]
              adaptation-rule pir adaptation-rule [cir adaptation-rule]
              burst-limit size [bytes | kilobytes]
              mbs mbs
              rate pir-rate [cir cir-rate]
              slope-policy hsm da-slope-policy-name
              wrr-weight weight
            exit
          exit

```

14.4.5 Configuring HSM DA Queue Group Overrides

When a queue group template is instantiated at the port, queue group queues can be overridden at the instance level using **hsm da-queue-overrides**.

The following output displays a hsm da-queue-group and overrides example configuration at the port level.

```

*A:Dut-A>config>port>ethernet# /configure port 4/1/1
*A:Dut-A>config>port# info
-----
    shutdown
  ethernet
    mode access
    egress
      exp-secondary-shaper "Exp_Shaper_1" create
      exit
    exit

```

```
access
  egress
    queue-group "100" instance 1 create
      hsmda-queue-override
        secondary-shaper "Exp_Shaper_1"
        wrr-policy "wrrRes"
        packet-byte-offset add 3
        queue 2 create
          rate 2222
        exit
      exit
    exit
  exit
exit
-----
```


14.5 Applying HSMDA Policies

HSMDA policies and values are associated in the following entities. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide* for command syntax and usage.

```

config
  — lag
    — port-type {standard | hsmda-ports}
    — no port-type
  — port
    — ethernet
      — access
      — egress
        — queue-group queue-group-name [create] [instance instance-id]
          — [no] hsmda-queue-override
            — packet-byte-offset {add add-bytes | subtract sub-bytes}
            — no packet-byte-offset
            — queue queue-id [create]
            — no queue queue-id
              — mbs {[0 to 2625] kilobytes | [0 to 2688000] bytes | default}
              — no mbs
              — rate pir-rate
              — no rate
              — slope-policy hsmda-slope-policy-name
              — no slope-policy
              — wrr-weight weight
              — no wrr-weight
            — secondary-shaper secondary-shaper-name
            — no secondary-shaper
            — wrr-policy wrr-policy-name
        — hsmda-scheduler-overrides
          — group group-id rate rate
          — no group
          — max-rate rate
          — no max-rate
          — scheduling-class class rate rate
          — scheduling-class class [weight weight-in-group]
          — no scheduling-class class

```

14.6 HSMDA Features

14.6.1 HSMDA LAG

The addition of the HSMDA affects Ethernet Link Aggregation Groups (LAGs). Due to the different behavior between a SAP created on a standard Ethernet MDA and a SAP created on an HSMDA, it is important to know the expected behavior at the time the SAP is created. A SAP can be created on a LAG that has no port members so at the time of creation the type of SAP may be unknown.

To remove the issue between SAP type and LAG port membership, the **config>lag>port-type {standard | hsmdda-ports}** command must be executed prior to adding any ports to the LAG. This command allows the type of ports that will be added to the LAG to be predefined. Without executing this command, HSMDA ports cannot be added to the LAG and, after execution, the LAG may only be populated with HSMDA ports.

The LAG port type restriction can only be changed prior to adding SAPs or binding network IP interfaces to the LAG. If the port type for the LAG must be changed, all ports, SAPs and IP interfaces must be removed from the LAG.

A LAG with HSMDA member ports cannot be configured as mode network.

A LAG with HSMDA member ports can only operate in link-level SLA distribution. Since hierarchical QoS is only supported for queues within a single queue group and a queue group is limited to a single egress port, it cannot spread a SAP's ingress or egress queue CIR and PIR parameters over multiple LAG links. Each SAP created on an HSMDA LAG is assigned a queue group for each link within the LAG. The CIR and PIR defined in the SAP ingress or SAP egress QoS policy is replicated for the queue ID in each queue group.

14.6.2 Billing

The HSMDA SAP and subscriber queue billing statistics collection process supports the same information as non-HSMDA objects with the addition of the exception counter information. Statistics from queues within a queue group that are not currently mapped to forwarding classes are removed. Since the queues always exist, the counter information for the unused queues will be presented by the underlying collection mechanisms. Because of a large amount of data that could potentially exist, longer statistics collection intervals can occur.

14.6.3 Resource Management

The HSMDA presents a different set of resource management features to the system than with non-HSMDAs. Since the HSMDA does not manage a pool of queues that are individually mapped to SAPs or subscribers, it cannot run out of queues on the XMA or MDA. Instead, a pool of available queue groups are managed and allocated on a per SAP or subscriber basis.

Each SAP and subscriber created is managed against a system-wide maximum. The maximum for SAPs and subscribers are each 64K within a single chassis.

14.6.4 HSMDA Queue Groups

A fundamental concept on an HSMDA is the queue group. Queue groups are not directly managed by the provisioner, they are indirectly assigned when creating SAPs or subscribers on the XMA and MDA. A queue group has eight queue members. The queues within the group are numbered from 1 through 8. When creating a SAP or subscriber associated with a port on the XMA or MDA, an ingress and egress queue group is allocated to the object. Every SAP and subscriber using the XMA or MDA has 8 ingress and 8 egress queues (whether they are in use or not). In the SAP ingress and egress QoS policies, the HSMDA queues within the group are represented by queue-id 1 through 8.

Each queue within the group has two RED slopes (managed by associating a slope policy to the queue), an MBS defined in bytes, a byte offset parameter used to add or subtract bytes to each packet handled by the queue for accounting purposes.

The queue group supports an aggregate shaper used to manage an aggregate rate limit for all queues within the group. Scheduling for queues within the group is stopped and started based on the rate set on the shaper.

Each queue group also supports 16 counter sets. Eight of the counters are the default counters used by packets assigned to each queue, respectively. The remaining eight are exception counters and are named Counter 1 through Counter 8.

The number of queue groups available is dependent on the HSMDA variant. The ESS variant supports up to 8K ingress and egress queue groups. The SR variant supports up to 20K ingress and egress queue groups. The ability to utilize all available queue groups is dependent on the type of IOM that is hosting the HSMDA.

14.6.4.1 Scheduling Classes

The HSMDA supports eight scheduler classes that are directly mapped to the queue-id (1 through 8) for each SAP and subscriber queue. The scheduler class is not an internal QoS policy-driven forwarding class. Forwarding classes within the system are used between the ingress and egress forwarding complexes and help the system to manage packet marking or remarking decisions and also are used to map each packet to an ingress and egress queue. It is possible to have two different QoS policies: the first that maps forwarding class AF (for example) to queue number 3 and the second may map AF to queue number 5. While the system will make certain common decisions based on the AF forwarding class, the fact that it is being mapped to different scheduler classes within the HSMDA will dictate that the scheduling of AF will be different for the two QoS policies based on the scheduler behavior.

14.6.4.2 Scheduling Class Weighted Groups

An HSMDA scheduler handles groups of queues based on each queue's identifier. All queues numbered 1 are automatically placed in scheduler class 1, queues numbered 2 are placed in scheduler class 2, and so on, up to queues numbered 8 being placed in scheduler class 8. Each scheduling class may be directly associated with a strict priority level or may be placed in one of two weighted groups. Each weighted group is used to map up to three scheduling classes into a single strict priority level and provides a weight for each member class of the group. Using the weighted groups allows for a mixture of strict and weighted scheduling between the scheduling classes. The scheduling classes mapped to a group must be consecutive in class order. If class 3 is placed into group 1, then the next class that may be placed into the group must be 2 or 4. If 4 were added to the group, then the next (and last) class that can be added to the group would be 5. If 2 had been added to the group instead of 4, then the next class would be limited to class 1.

14.6.4.3 Scheduler Strict Priority Levels

The scheduler maintains 8 strict levels with strict level 8 being the highest priority and strict level 1 being the lowest. For each strict priority level, either the scheduler class with the same ID (1 through 8) may be mapped to the strict level, or one of the two weighted groups may be mapped to the strict level. If a scheduling class is not mapped to a weighted group, the class is instead mapped to its relative strict

scheduling level. When a scheduling class is mapped to a weighted group, it is removed from the strict level and shares the strict level assigned to the weighted group. The weighted group itself is mapped to the inherent strict scheduling level of the highest member scheduling class. Therefore, when weighted scheduling class groups are used, fewer strict levels are active on the scheduler.

14.6.4.4 Strict Priority Level PIR

The scheduler supports a strict scheduling level PIR that limits the amount of bandwidth allowed for the level. The rate is defined in increments of megabits per second and may be set to max (the default setting), which disables the shaping function. The shaping rate is not defined on the strict priority level, but is inherited from the scheduling class or weighted group that is mapped to the strict level. The scheduler includes the full Ethernet frame encapsulation overhead when updating the priority level PIR, including the 12 byte inter-frame gap and the 8 byte preamble.

14.6.4.5 Scheduler Maximum Rate

A maximum scheduling rate may be defined for the scheduler. The rate is specified in megabits per second and the default rate is max, which allows the scheduler to operate without a set limit. When the HSMDA scheduling policy is applied to an egress port, the maximum scheduling rate may be used to define a rate less than the available line rate of the port. When the HSMDA scheduler policy is applied to the ingress path of an HSMDA, it sets the maximum ingress bandwidth for the XMA or MDA (not usually a desirable action). The scheduler includes the full Ethernet frame encapsulation overhead when updating the scheduler level PIR, including the 12 byte inter-frame gap and the 8 byte preamble.

14.6.4.6 HSMDA Scheduler Policy Overrides

When an HSMDA scheduler is applied to an egress port or to an ingress HSMDA, the various parameters may be overridden. This allows an HSMDA scheduler policy to be adapted to changing needs on a port or HSMDA basis without requiring a new policy to be created.

14.6.4.7 Orphan Queues

Unlike port- or service-based virtual scheduling behavior, the HSMDA schedulers do not need to deal with orphaned queues (queues without an explicit scheduler parent defined). Every queue on an HSMDA is implicitly mapped to the scheduler based on the queues identifier.

14.6.4.8 Default HSMDA Scheduling Policy

An HSMDA scheduling policy with the name default always exists on the system and does not need to be created. The default policy cannot be modified or deleted. Attempting to delete the default policy using the **no hsmda-scheduler-policy default** command will return an error without changing the default policy.

The default policy contains the parameters listed in [Table 114](#).

Table 114 HSMDA Scheduling Policy Default Values

Command	Default
description	no description
max-rate	no max-rate
group	group 1 rate max
	group 2 rate max
scheduling-class	scheduling-class 1 rate max
	scheduling-class 2 rate max
	scheduling-class 3 rate max
	scheduling-class 4 rate max
	scheduling-class 5 rate max
	scheduling-class 6 rate max
	scheduling-class 7 rate max
	scheduling-class 8 rate max

14.7 HSMDA Command Reference

14.7.1 Command Hierarchies

```

config
  — qos
    — hsmda-pool-policy policy-name [create]
    — no hsmda-pool-policy policy-name
      — class-tier
        — class-pool class-pool-id root-parent root-pool-id allocation-percent
          percent-of-parent-pool
        — no class-pool class-pool-id
      — description description-string
      — no description
      — root-tier
        — root-pool root-pool-id allocation-weight pool-weight
        — no root-pool root-pool-id
      — system-reserve percentage-of-buffers
      — no system-reserve
    — hsmda-scheduler-policy scheduler-policy-name [create]
    — no hsmda-scheduler-policy scheduler-policy-name
      — description description-string
      — no description
      — group group-id rate rate [burst-limit size] [bytes | kilobytes]
      — no group group-id
      — max-rate rate [burst-limit size] [bytes | kilobytes]
      — no max-rate
      — scheduling-class class group group-id [weight weight-in-group]
      — scheduling-class class rate rate [burst-limit size] [bytes | kilobytes]
      — no scheduling-class class
    — hsmda-slope-policy policy-name [create]
    — no hsmda-slope-policy policy-name
      — description description-string
      — no description
      — high-slope
        — max-depth percent-of-queue-depth
        — no max-depth
        — max-prob discard-probability-in-percent
        — no max-prob
        — [no] shutdown
        — start-depth percent-of-queue-depth
        — no start-depth
      — low-slope
        — max-depth percent-of-queue-depth
        — no max-depth
        — max-prob discard-probability-in-percent
        — no max-prob
        — [no] shutdown
        — start-depth percent-of-queue-depth

```

- **no start-depth**
- **hsmda-wrr-policy** *policy-name* [**create**]
- **no hsmda-wrr-policy** *policy-name*
 - **class-agg-weight** *weight*
 - **no class-agg-weight**
 - **description** *description-string*
 - **no description**
 - **include-queues** *queues*
 - **no include-queues**
 - **schedule-using-class** *class*
 - **no schedule-using-class**

14.7.1.1 Show Commands

- show**
- **qos**
 - **hsmda-pool-policy** [*hsmda-pool-policy-name*] [**associations**] [**detail**]
 - **hsmda-pools** **mda** *mda-id* {**ingress** | **egress**} [**detail**]
 - **hsmda-pools** **port** *port-id* **class** *class-id*
 - **hsmda-pools** **port** *port-id* **queue-group** *queue-group-name* **instance** *instance-id*
 - **hsmda-pools** **sap** *sap-id* [**port** *port-id*]
 - **hsmda-pools** [**port** *port-id*] **subscriber** *sub-ident*
 - **hsmda-pools** **port** *port-id* **network-queues**
 - **hsmda-scheduler-hierarchy** **port** *port-id* [**interval** *time*]
 - **hsmda-scheduler-hierarchy** **port** *port-id* {**shapers** | **shaper** *shaper-name*} [**interval** *time*] [**detail**]
 - **hsmda-scheduler-hierarchy** **mda** *mda-id* [**interval** *time*]
 - **hsmda-scheduler-hierarchy** **sap** *sap-id* [**ingress** | **egress**] [**interval** *time*]
 - **hsmda-scheduler-hierarchy** **subscriber** *sub-id* [**ingress** | **egress**] [**interval** *time*]
 - **hsmda-scheduler-policy** [*hsmda-scheduler-policy-name*] [**associations**] [**detail**]
 - **hsmda-slope-policy** [*hsmda-slope-policy-name*] [**associations**] [**detail**]

14.7.2 Command Descriptions

14.7.2.1 Configuration Commands

14.7.2.1.1 Generic Commands

description

Syntax **description** *description-string*

	no description
Context	config>qos>hsmda-pool-policy config>qos>hsmda-scheduler-policy config>qos>hsmda-slope-policy config>qos>hsmda-wrr-policy
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	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

shutdown

Syntax	[no] shutdown
Context	config>qos>hsmda-slope-policy>high-slope config>qos>hsmda-slope-policy>low-slope
Description	This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted. Services are created in the administratively down (shutdown) state. When a no shutdown command is entered, the service becomes administratively up, then tries to enter the operationally up state. The no form of this command places the entity into an administratively enabled state.

14.7.2.1.2 HSMda Pool QoS Policy Commands

hsmda-pool-policy

Syntax	hsmda-pool-policy <i>policy-name</i> [create] no hsmda-pool-policy <i>policy-name</i>
Context	config>qos

Description	This command enters the context to configure a HSM DA pool policy parameters. Each policy must be uniquely named within the system. Names of up to 32 ASCII characters composed of printable, 7-bit ASCII characters are supported. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. A policy must be defined prior to applying the policy name to an HSM DA entity. The no form of this command removes the specified HSM DA pool policy from the configuration. If the HSM DA pool policy is associated with an HSM DA, the command will fail.
Parameters	<i>policy-name</i> — Specifies the name of the pool policy up to 32 characters. create — The create keyword is required when the create command is enabled in the environment context to create a new HSM DA scheduler policy. This keyword is not required when the protection is disabled. The keyword is ignored when the HSM DA scheduler policy already exists.

class-tier

Syntax	class-tier
Context	config>qos>hsm da-pool-policy
Description	This class enables the context to configure class pool tier parameters. Within the class-tier context, class pools may be associated with a root pool and are sized as a percentage of the root pool's size.

class-pool

Syntax	class-pool <i>class-pool-id</i> root-parent <i>root-pool-id</i> allocation-percent <i>percent-of-parent-pool</i> no class-pool <i>class-pool-id</i>
Context	config>qos>hsm da-pool-policy
Description	This command specifies a class pool's root pool parent and define the buffer allocation percentage used for sizing the class pool relative to the parent pool's size. Eight class pools exist and do not need to be created. Class pools function as a scheduling class aggregate buffer control mechanism and define the total number of buffers a given scheduling class may consume for all ports. Each class pool must be placed in a root pool hierarchy. This is accomplished by the root-parent keyword and root-pool-id parameter. A class pool cannot be parented by a root pool that currently has an allocation-weight parameter set to 0. When a class pool is parented by a root pool, that root pool will not allow the allocation-weight to be set to 0.

The allocation-percent keyword and associated percent-of-parent-pool parameter indirectly specifies the size of the class pool. The percent value is multiplied by the size of the root pool to derive the class pool size. If the percent value is changed, or the size of the root pool changes, the class pool and port class pools associated with the class pool must be resized. The sum percent values for the class pools associated with a root pool may exceed 100%, allowing the class pools to oversubscribe the root pool. In this case, it is possible for the class pool to indicate that buffers remain when the root pool has exhausted its available buffers.

When queues associated with the class pool’s scheduling class request buffers due to packet arrival, the port class pool, class pool, and root pool must all have sufficient buffers available to place the packet into the queue. If sufficient buffers are not available, the packet will be discarded by the queue.

The **no** form of this command restores the default root-parent and allocation-percent value for a class pool. Based on the class pool, the restored default values may differ.

Parameters *class-pool-id* — Specifies which class pool is being modified. This parameter is a required when executing the class-pool command.

Values 1 to 8

root-pool-id — Specifies the parent root pool to which the class pool will be associated. All class pool parent associations are output when save config or show config is executed, regardless of whether the default value is currently set. The root-parent keyword is required and must precede the root-pool-id parameter.

Values 1 to 8

[Table 115](#) lists the root pool class types and their associated IDs.

Table 115 Root Pool ID Class Pool

Unit	Integer
Default class-pool 1	1
Default class-pool 2	1
Default class-pool 3	1
Default class-pool 4	1
Default class-pool 5	1
Default class-pool 6	2
Default class-pool 7	2
Default class-pool 8	2

percent-of-parent-pool — Defines the percentage of the root pool's size to define the size of the class pool. The value is specified as a percentage with two decimal places (100th of a percent). All class pool percentage values are output when save config or show config is executed, regardless of whether the default value is currently set. The allocation-percent keyword is required and must precede the percent-of-parent-pool parameter.

root-tier

Syntax	root-tier
Context	config>qos>hsmda-pool-policy
Description	This command enables the root pool tier context of an HSMDA pool policy. Within the root-tier context, root pools may be sized by defining each root pool's weight.

root-pool

Syntax	root-pool <i>root-pool-id</i> allocation-weight <i>pool-weight</i> no root-pool <i>root-pool-id</i>
Context	config>qos>hsmda-pool-policy
Description	This command defines the buffer allocation weight for a specific root pool. Eight root pools exist and do not need to be created. The allocation-weight parameter is used to specify the weight that will be applied to the pool and is divided by the sum of all root pool weights to derive the pool's buffer allocation factor. The amount of buffers remaining after the system-reserve percentage is applied is multiplied by the buffer allocation factor to derive the pool size.

Root pools function as an oversubscription control mechanism. A root pool acts as the root of a hierarchy of buffer pools and queues with respect to buffer allocation. Since the sum of the root pool sizes will not exceed the total number of buffers available, the number of buffers indicated by the root pools size will always be available to the queues within the root pools hierarchy, queues from one hierarchy can never steal buffers from another.

A root pool hierarchy is based on the dynamic parenting of a root pool to one or more class pools. A class pool represents the buffering allowed for a given scheduling class. Each class pool is sized as a percentage of the root pool to which it is parented. The sum of the class pools percentages for a root pool may be greater than 100 percent, which allows the root pool to be oversubscribed. This may be beneficial when large fluctuations in class-based buffer utilization are expected and a given class should be allowed to exceed its fair share of buffering.

Port queues are tied to root pools through the scheduling class of the queue (indicated by the queue-ID). A queue on scheduling class 3 will be mapped to class pool 3 and indirectly tied to the root pool associated with class pool 3.

A root pool with an allocation-weight set to 0 is considered inactive and will not be allocated buffers. Class pools cannot be parented to a root pool with a weight equal set to 0. When a class pool is associated with a root pool, the root pool's weight cannot be set to 0.

When a root pool's allocation weight is modified, all root pools, class pools, and port class pool sizes are reevaluated and modified when necessary.

The **no** form of this command restores the default allocation-weight value to a root pool. Root pool 1 has a different default weight than root pools 2 through 8. The no root-pool command will fail for root pools 2 through 8 if the root pool is currently parented to a class pool.

- Parameters

root-pool-id — Specifies the root pool ID value. This value is a required parameter when executing the root-pool command and specifies which root pool is being modified.

Values 1 to 8

Default none

pool-weight — The pool-weight parameter defines the weight of the specified root-pool-id and is used by the system to calculate the size of the root buffer pool. The allocation-weight keyword is required and must precede the pool-weight parameter. Setting pool-weight to 0 disables the pool and prevents the root pool from being a parent to any class pools. Root pool 1 cannot be set with an allocation weight of 0.

[Table 116](#) lists the pool weight values and their associated IDs.

Table 116 Pool Weight Values

Unit:	Integer
Range root-pool 1	1 to 100
Range root-pool 2	0 to 100
Range root-pool 3	0 to 100
Range root-pool 4	0 to 100
Range root-pool 5	0 to 100
Range root-pool 6	0 to 100
Range root-pool 7	0 to 100
Range root-pool 8	0 to 100
Default root-pool 1	75
Default root-pool 2	25
Default root-pool 3	0
Default root-pool 4	0
Default root-pool 5	0

Table 116 Pool Weight Values (Continued)

Unit:	Integer
Default root-pool 6	0
Default root-pool 7	0
Default root-pool 8	0

system-reserve

Syntax **system-reserve** *percentage-of-buffers*
no system-reserve

Context config>qos>hsmdda-pool-policy

Description This command defines the amount of HSMDDA buffers that will be set aside for internal system use. By default, 10% of the total buffer space is reserved for system internal queues. The command is provided for the case where the reserved buffer space is either insufficient or excessive. Use care when modifying this value. When the system reserve value is changed, all the provisioned port class, class, and root pool sizes will be reevaluated and possibly changed.

The **no** form of this command restores the default system reserve value.

Parameters *percentage-of-buffers* — Specifies the system reserve value.

Values 1.00 to 30.00

Default 10.00

14.7.2.1.3 HSMDDA Scheduler QoS Policy Commands

hsmdda-scheduler-policy

Syntax **hsmdda-scheduler-policy** *scheduler-policy-name* [**create**]
no hsmdda-scheduler-policy *scheduler-policy-name*

Context config>qos

Description This command configures HSMMDA scheduler policy parameters. HSMMDA scheduler policies can be assigned to an egress HSMMDA port or as the ingress control scheduler between the HSMMDA and the ingress forwarding plane. The policy contains commands to provision the scheduling behavior of a set of HSMMDA scheduler classes. When assigned to an HSMMDA egress port, the policy defines the scheduling behavior for all queues associated with the egress port. When assigned to the ingress path of an HSMMDA, the policy defines the scheduling behavior for all ingress queues on the HSMMDA (regardless of ingress port).

Up to 1000 HSMMDA scheduler policies can be configured per system.

The **no** form of this command removes an HSMMDA scheduler policy from the system. If the policy is associated with an egress port or ingress HSMMDA, the command will fail.

Parameters *scheduler-policy-name* — Each HSMMDA scheduler policy must be uniquely named within the system. Names of up to 32 ASCII characters composed of printable, 7-bit ASCII characters are supported. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. An HSMMDA scheduler policy must be defined prior to applying the policy name to an HSMMDA egress port or to an ingress HSMMDA.

create — The **create** keyword is required when the **create** command is enabled in the **environment** context to create a new HSMMDA scheduler policy. This keyword is not required when the protection is disabled. The keyword is ignored when the HSMMDA scheduler policy already exists.

group

Syntax **group** *group-id* **rate** *rate* [**burst-limit** *size*] [**bytes** | **kilobytes**]
no group *group-id*

Context config>qos>hsmmda-scheduler-policy

Description This command defines the maximum rate allowed for the scheduling classes mapped to the specified *group-id*. A group is a scheduling entity used to combine up to three consecutive scheduling classes into a single strict priority level. Each scheduling class within the group has an associated weight. When the scheduler is operating at the strict level associated with the group, the ratio of bandwidth allocated to each scheduling class within the group during congestion at that strict level is relative to the ratio of the weight of each member. The bandwidth is allocated in a work-conserving fashion and is sensitive to packet size up to the maximum rate defined for the group.

The **no** form of this command reverts the specified weighted scheduling class group rate limit to the default setting.

Parameters *group-id* — Defines the maximum rate allowed for the scheduling classes mapped to the specified group-id.

Values 1, 2

rate — Specifies the maximum rate in megabits per second. If the **max** keyword is used with the **rate** parameter, the bandwidth limitation is removed from the group.

Values 1 to 40000000, max

Default max

size — Specifies the burst limit size.

Values 1 to 327680

bytes — Specifies to use the size in bytes.

kilobytes — Specifies to use the size in kilobytes.

max-rate

Syntax **max-rate** *rate* [**burst-limit** *size*] [**bytes** | **kilobytes**]

no max-rate

Context config>qos>hsmda-scheduler-policy

Description This command defines an explicit maximum frame-based bandwidth limit for the HSMDA scheduler policy scheduler. If a **max-rate** is defined that is smaller than the port rate, the port will be rate limited to the expressed megabits-per-second value. This command should be used with caution if the policy may be applied on ingress for an HSMDA as the total port ingress rate will be limited to the defined maximum rate.

This command can be executed at any time for any non-default existing HSMDA scheduler policy. When a new maximum 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 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 maximum rate value defined in the policy can be overridden on each scheduler instance. If the maximum rate is explicitly defined as an override on a port or ingress HSMDA, the policy's max-rate value has no effect.

The **no** form of this command removes an explicit rate value from the HSMDA scheduler policy. When removed, all instances of the scheduler policy on egress ports or ingress HSMDAs are allowed to run at the available line rate unless the instance has a max-rate override in place.

Parameters *rate* — Specifies an explicit maximum frame-based bandwidth limit for the HSMDA scheduler policy scheduler.

Values 1 to 40000000 megabits per second, max

size — Specifies the burst limit size.

Values 1 to 327680

bytes — Specifies to use the size in bytes.

kilobytes — Specifies to use the size in kilobytes.

scheduling-class

Syntax	scheduling-class <i>class</i> group <i>group-id</i> [weight <i>weight-in-group</i>] scheduling-class <i>class</i> rate <i>rate</i> [burst-limit <i>size</i>] [bytes kilobytes] no scheduling-class <i>class</i>
Context	config>qos>hsmdda-scheduler-policy
Description	This command configures the behavior of a specific scheduling class on all HSMDDA schedulers associated with the policy. The scheduling-class command performs one of two operations: configure a maximum rate for the scheduling class or place the scheduling class into one of the two available weighted scheduling groups. The two operations are mutually exclusive.

By default, none of the scheduling classes are members of either weighted scheduling group and each class is set to a rate limit of **max** (meaning that no rate limit is applied).

Specifying Scheduling Class Rate (or Removing Class from Group)

If the **scheduling-class** command is executed with the **rate** keyword specified, either **max** or a specified megabits-per-second rate must follow. If class-id had previously been mapped into one of the two weighted scheduling groups, the class will be removed. However, if removing the class from the group will cause the group to no longer have contiguous class members, the command will fail with no effect on the specified class. A non-contiguous grouping error will be returned, specifying the weighted group. The lowest or highest members within a weighted group must be removed prior to removing the middle member. For example, if scheduling classes 3, 4, and 5 were members of weighted group 1, class 4 cannot be removed first.

The **scheduling-class** command using the **rate** keyword will also fail in the event that an override for the group weight is in place on the scheduling class within a scheduler associated with the policy. The override command is expecting the class to be associated with a weighted scheduling group and the policy rate definition is attempting to remove the class from the group. An override mismatch error will be generated, specifying the scheduling object where the override exists (SAP, subscriber, or ingress HSMDDA).

When a rate has been successfully defined for a scheduling class, the specified rate is automatically updated on all HSMDDA scheduler instances associated with the scheduling policy. The exception is where the scheduler instance has a local override for the rate on the scheduling class.

Specifying Scheduling Class Weighted Group Membership

If the scheduling-class command is executed with the group keyword specified, group-id must follow. Two weighted scheduling groups are allowed, numbered 1 and 2. Along with the group, the weight keyword is used to specify the weight the scheduling class within the group. If weight is not specified, the default weight of 1 will be used. Similar to the rate action of the command, the group version will fail if the scheduling class ID is not consecutive with the class members that are currently members of the weighted scheduling group. The command will have no effect on the current scheduling class settings and a non-contiguous grouping error will be returned, specifying the weighted scheduling group and the current group members.

The scheduling-class command will also fail using the group keyword when a rate override for the scheduling class exists on an HSM DA scheduler instance associated with the policy. The rate override for the scheduling class indicates the class is directly attached to a strict priority level, conflicting with the policy group keyword trying to place the class in the specified group. The command will fail without effecting the scheduling class definition on the policy and return an override-mismatch error specifying the scheduling object where the override exists.

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.

Other Override Constraints

The scheduling overrides cannot change or remove a scheduling class from a policy-defined weighted group membership.

The **no** form of this command returns the scheduling class represented by class-id to the default behavior. The default behavior for a scheduling class is to not be a member of either weighted scheduling class groups and have a rate set to **max**. The **no** form of this command will fail if the scheduling class is currently a member of one of the weighted scheduling class groups and a weight override is in effect on a scheduling object for the class. An override mismatch error will be returned, specifying the scheduling object where the override exists.

Parameters *class* — specifies the weight the QMDA port scheduler policy should apply to this policy level within the group it belongs to.

Values 1 to 8

group-id — If the **scheduling-class** command is executed with the **group** keyword specified, a *group-id* must be specified. Two weighted scheduling groups are allowed, numbered 1 and 2. With the **group**, the **weight** keyword specifies the weight the scheduling class within the group. If weight is not specified, the default weight is 1. Similar to the rate action of the command, the group version will fail if the scheduling class ID is not consecutive with the class members currently members of the weighted scheduling group. The command will have no effect on the current scheduling class settings and a non-contiguous grouping error will be returned, specifying the weighted scheduling group and the current group members.

The **scheduling-class** command will also fail using the **group** keyword when a **rate** override for the scheduling class exists on an HSMDA scheduler instance associated with the policy. The rate override for the scheduling class indicates the class is directly attached to a strict priority level, conflicting with the policy group keyword trying to place the class in the specified group. The command will fail without effecting the scheduling class definition on the policy and return an override-mismatch error specifying the scheduling object where the override exists.

The configured priority level rate limits can 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.

Values 1, 2

weight-in-group — Specifies the relative weight of *class-id* to the other scheduling classes within the group. If group is specified without the weight keyword, a default weight is 1.

Values 1 to 100

rate — If the **scheduling-class** command is executed with the **rate** keyword specified, either **max** or a specified megabits-per-second rate must follow. If class-id was previously mapped into one of the two weighted scheduling groups, the class will be removed. However, if removing the class from the group will cause the group to no longer have contiguous class members, the command will fail with no effect on the specified class. A non-contiguous grouping error will be returned, specifying the weighted group. The lowest or highest members within a weighted group must be removed prior to removing the middle member. For example, if scheduling classes 3, 4 and 5 were members of weighted group 1, class 4 cannot be removed first.

The **scheduling-class** command using the **rate** keyword will also fail in the event that an override for the group weight is in place on the scheduling class within a scheduler associated with the policy. The override command is expecting the class to be associated with a weighted scheduling group and the policy rate definition is attempting to remove the class from the group. An override mismatch error will be generated, specifying the scheduling object where the override exists (such as a SAP, subscriber, or ingress HSMDA).

When a rate has been successfully defined for a scheduling class, the specified rate is automatically updated on all HSMMDA scheduler instances associated with the scheduling policy. The exception is where the scheduler instance has a local override for the rate on the scheduling class.

The **max** keyword specifies that a limit is not enforced for the specified *class-id* and that the *class-id* is not a member of a weighted scheduling class group. The **max** keyword is mutually exclusive to the kilobits-per-second parameter and when specified, must directly follow the rate keyword. Setting the rate of the class to **max** will fail when the class currently has a group weight override defined on a scheduling object (SAP, subscriber profile, or ingress HSMMDA).

Values 1 to 40000000, max

14.7.2.1.4 HSMMDA Slope QoS Policy Commands

hsmmda-slope-policy

Syntax	hsmmda-slope-policy <i>policy-name</i> [create] no hsmmda-slope-policy <i>policy-name</i>
Context	config>qos
Description	This command creates an HSMMDA RED slope policy. The policy may be assigned to an ingress or egress HSMMDA queue. The policy contains the Maximum Buffer Size (MBS) that will be applied to the queue and the high- and low-priority RED slope definitions. The function of the MBS and RED slopes is to provide congestion control for an HSMMDA queue. The MBS parameter defines the maximum depth a queue may reach when accepting packets. The low- and high-priority RED slopes provide for random early detection of congestion and slope-based discards based on queue depth.

An HSMMDA slope policy can be applied to queues defined in the SAP ingress and SAP egress QoS policy HSMMDA queues context. When an HSMMDA slope policy is applied to a SAP QoS policy queue, it cannot be deleted. Any edits to the policy are updated to all HSMMDA queues indirectly associated with the policy.

Default HSMMDA Slope Policy

An HSMMDA slope policy named "default" always exists on the system and does not need to be created. The default policy is automatically applied to all HSMMDA queues unless another HSMMDA slope policy is specified for the queue. The default policy cannot be modified or deleted. Attempting to execute the **no hsmmda-slope-policy default** command results in an error.

The **no** form of this command removes the specified HSMMDA slope policy from the configuration. If the HSMMDA slope policy is currently associated with an HSMMDA queue, the command will fail.

Table 117 lists the default slope policy parameters.

Table 117 HSMDB Default Slope Policy Values

Parameter		Default Value
queue-mbs		16,800 bytes
high-slope		
	start-depth	100.00
	max-depth	100.00
	max-prob	100.00
	shutdown	no shutdown
high-slope		
	start-depth	90.00
	max-depth	90.00
	max-prob	100.00
	shutdown	no shutdown

- Parameters** *policy-name* — Specifies a HSMDB slope policy. Each HSMDB slope policy must be uniquely named within the system. Names of up to 32 ASCII characters composed of printable, 7-bit ASCII characters are supported. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. An HSMDB slope policy must exist prior to applying the policy name to an HSMDB queue; 1024 policies are allowed per system.
- create** — The **create** keyword is required when the **create** command is enabled in the **environment** context to create a new HSMDB scheduler policy. This keyword is not required when the protection is disabled. The keyword is ignored when the HSMDB scheduler policy already exists.

high-slope

- Syntax** **high-slope**
- Context** config>qos>hsmdb-slope-policy
- Description** This command enables the high-priority RED slope context of an HSMDB slope policy. Within the high-slope context, the high-priority RED slope configuration commands defining the start of slope, end of slope, and maximum probability points may be executed.

For ingress, packets classified as priority high or profile in are mapped to the high-priority RED slope for queue congestion management.

At egress, packets received from ingress as in-profile, or that are reclassified to inplus-profile or in-profile at egress, are mapped to the high-priority RED slope for queue congestion management. In-profile is derived at ingress either from within-CIR profiling or from explicit profile in classification.

max-depth

Syntax	max-depth <i>percent-of-queue-depth</i> no max-depth
Context	config>qos>hsmda-slope-policy>high-slope config>qos>hsmda-slope-policy>low-slope
Description	This command defines the ending queue depth point for a slope relative to the maximum depth of the queue-mbs value of the HSMDBA slope policy. At the point the slope ends, it has risen from the starting queue depth value with a discard probability of zero. At max-depth, the slope will have risen to the discard probability defined by max-depth at which point the slope rises directly to a discard probability of 100%. If the queue depth has reached the point defined by max-depth, all packets associated with the slope will be discarded. The defined percent-of-queue-depth for max-depth is defined as a percentage of the queue-mbs value. The value defined as the maximum depth must be greater than or equal to the current percentage value for start-depth. If the defined value is less than start-depth, the max-depth command will fail with no change to the current value. If the start-depth value is greater than the desired max-depth value, first change start-depth to a value equal to or less than the desired max-depth. The no form of this command restores the default ending point percentage value for the slope. The low and high slopes have different default values. If the default value is less than the current start-depth value, the no max-depth command will fail.
Parameters	<i>percent-of-queue-depth</i> — Specifies the max depth for the high or low slopes The percent-of-queue-depth parameter is required when executing the max-depth command. It is expressed as a percentage value with two decimal places (100th of a percent) accuracy. Specifying the fractional portion is optional.

max-prob

Syntax	max-prob <i>discard-probability-in-percent</i> no max-prob
Context	config>qos>hsmda-slope-policy>high-slope

config>qos>hsmda-slope-policy>low-slope

Description This command defines the slopes maximum probability point where the slope ends and the discard probability rises directly to 100%. Together with the max-depth command, the max-prob value defines the end of the slope.

The defined percent-of-discard-probability for max-prob is defined as a percentage of a 100% discard probability. If a value of 75% is defined, near the end of the slope close to 75% of the packets will be discarded. At the end of the slope, the discard probability rises to 100% and all packets are discarded.

Parameters *discard-probability-in-percent* — Specifies the max probability percentage for the high or low slopes. The percent-of-discard-probability parameter is required when executing the max-prob command. It is expressed as a percentage value with two decimal places (100th of a percent) accuracy. Specifying the fractional portion is optional.

start-depth

Syntax **start-depth** *percent-of-queue-depth*
no start-depth

Context config>qos>hsmda-slope-policy>high-slope
config>qos>hsmda-slope-policy>low-slope

Description This command defines the starting point for a slope relative to the maximum depth of the queue-mbs value of the HSMDA slope policy. At the point the slope starts, it rises from a discard probability of zero until it reaches the max-depth and max-prob points defining where the slope rises directly to a discard probability of 100%.

As packets arrive at the queue, a random number is generated that ranges from 0 to 100% discard probability. The packet will be associated with either the high-priority slope or the low-priority slope. The current queue depth plots the position where the slope will be evaluated. If the random number is equal to or greater than the slope's discard probability at the current depth, the packet is eligible to enter the queue. If the random number is less than the slope discard probability, the packet is discarded.

The defined *percent-of-queue-depth* value for the start-depth command is defined as a percentage of the **queue-mbs** value. The value defined for the start depth must be less than or equal to the current percentage value for max-depth. If the defined value is greater than **max-depth**, the **start-depth** command will fail with no change to the current value. If the **max-depth** value is less than the desired **start-depth** value, first change **max-depth** to a value equal to or greater than the desired **start-depth**.

The **no** form of this command restores the default start point percentage value for the slope. The low and high slopes have different default values. If the default value is greater than the current max-depth value, the **no start-depth** command will fail.

Parameters *percent-of-queue-depth* — Specifies the start depth for the high or low slopes.

low-slope

Syntax	low-slope
Context	config>qos>hsmda-slope-policy
Description	This command enables the low-priority RED slope context of an HSMDA slope policy. Within the low-slope context, the low-priority RED slope configuration commands defining the start of slope, end of slope, and maximum probability points may be executed. For ingress, packets classified as low priority or out-of-profile are mapped to the low priority RED slope for queue congestion management. At egress, packets received from ingress as out-of-profile, or that are reclassified to out-profile or exceed-profile at egress, are mapped to the low-priority RED slope for queue congestion management. Out-of-profile is derived at ingress either from above-CIR profiling or from explicit profile out classification.

14.7.2.1.5 HSMDA WRR QoS Policy Commands

hsmda-wrr-policy

Syntax	hsmda-wrr-policy <i>policy-name</i> [create] no hsmda-wrr-policy <i>policy-name</i>
Context	config>qos
Description	This command creates an HSMDA weighted-round-robin (WRR) scheduling loop policy. The policy may be assigned to an egress HSMDA queue. The no form of this command removes the specified HSMDA WRR policy from the configuration. If the HSMDA WRR policy is currently associated with an HSMDA queue, the command will fail.
Parameters	<i>policy-name</i> — Specifies a HSMDA WRR policy. Each HSMDA WRR policy must be uniquely named within the system. Names of up to 32 ASCII characters composed of printable, 7-bit ASCII characters are supported. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. An HSMDA WRR policy must exist prior to applying the policy name to an HSMDA queue. Up to 32 WRR policies are supported on the system. create — The create keyword is required when the create command is enabled in the environment context to create a new HSMDA WRR policy. This keyword is not required when the protection is disabled. The keyword is ignored when the HSMDA WRR policy already exists.

class-agg-weight

Syntax	class-agg-weight <i>weight</i> no class-agg-weight
Context	config>qos>hsmda-wrr-policy
Description	This command specifies the aggregate weight within the scheduling class for the HSM DA WRR scheduling loop policy. The no form of this command removes the explicit aggregate weight value from the HSM DA WRR policy.
Parameters	<i>weight</i> — Specifies the aggregate weight for the HSM DA WRR policy. Values 1, 2, 4, 8

include-queues

Syntax	include-queues <i>queues</i> no include-queues
Context	config>qos>hsmda-wrr-policy
Description	This command includes either the first two (1 and 2) queues or the first three (1, 2, and 3) queues into the HSM DA WRR scheduling loop policy. This command defines how many queues are members of the scheduling loop. The no form of this command removes the explicit queues from the HSM DA WRR policy.
Parameters	<i>queues</i> — Specifies which queues to include. 1-2 specifies that the first two (1 and 2) queues are included in the HSM DA WRR policy. 1-3 specifies that the first three (1, 2, and 3) queues are included in the HSM DA WRR policy. Values 1-2, 1-3

schedule-using-class

Syntax	scheduling-using-class <i>class</i> no scheduling-using-class
Context	config>qos>hsmda-wrr-policy
Description	This command specifies the scheduling class for the HSM DA WRR scheduling loop policy. This command specifies which scheduling class the scheduling loop will participate with on the system. The no form of this command removes the explicit scheduling class value from the HSM DA WRR policy.


```

2          2          100.00
3          3          100.00
4          4          100.00
5          5          100.00
6          6          100.00
7          7          100.00
8          8          100.00
=====

```

```

=====
Root Tier
=====

```

```

Root Pool      Root Weight
-----

```

```

1          100
2          100
3          100
4          100
5          100
6          100
7          100
8          100
=====

```

```

-----
Associations
-----

```

```

- MDA Ingress: 4/1

```

```

=====
*A:Dut-A#

```

```

*A:Dut-A# show qos hsmda-pool-policy ingPoolPol association

```

```

=====
Qos HSMda Pool Policy
=====

```

```

Policy Name   : ingPoolPol

```

```

Description   : (Not Specified)

```

```

-----
Associations
-----

```

```

- MDA Ingress: 4/1

```

```

=====
*A:Dut-A#

```

Table 118 **Show HSMda Pool Policy Output Fields**

Label	Description
Policy Name	Displays the policy name.
Description	Displays the policy description.

Table 118 Show HSMDB Pool Policy Output Fields (Continued)

Label	Description (Continued)
Sys. Reserve	Displays the amount of HSMDB buffers that will be set aside for internal system use.
Class Pool	
Root Parent	
Alloc. Percent	
Abv Off PIR Prcnt	
Abv Off Rate	
BWD Gran PIR Prcnt	Displays the percentage of the child's administrative PIR that is used as the rounding step value. If a value of 0 or 0.00 is used, the system will interpret this equivalent to no granularity .
BWD Gran Rate	Displays the rate, in kb/s, that is used as the rounding step value. If a value of 0 is used, the system will interpret this equivalent to no granularity .
Fast Start	Displays whether fast detection of initial bandwidth on a child policer or queue associated with the policy is enabled.
Fast Stop	Displays whether fast detection of lack of offered rate on a child policer or queue associated with the policy is enabled.
Time Average Factor	Displays the time average factor used to calculate the weight of the new offered rate using a portion of the previous offered rate. If dec-only is true, the time average factor adjustment is only applied if the new offered rate is decreasing compared to the previous offered rate.
Hi Rate Hold Time	Displays the time period that the current offered rate is maintained for a child policer or queue when it is seen that the offered rate is decreasing. The offered measurement that triggers the hold time is used when the hold timer expires, unless a higher offered rate is seen in the interim. When a higher rate is observed, the hold timer is canceled and the higher offered rate is used immediately. If active-min-only is true, the respective rate or percentage is treated as the minimum offered rate for a queue, only when the queue has an actual non-zero offered rate.
Sample Interval	Displays the number of intervening sample periods before a new offered rate is measured.
Enqueue on PIR Zero	

Table 118 Show HSMDB Pool Policy Output Fields (Continued)

Label	Description (Continued)
Limit PIR Zero Drain	If true, the system to use the minimum configurable PIR instead of an HQoS derived zero operational PIR.
Init LUB with min PIR	If true, new queues associated with a LUB context are assigned a minimum rate operational PIR to be applied to the queue for use by enqueued packets prior to an HQoS iteration.
Internal Sched Wght Mode	Specifies the internal scheduler weight: default — Specifies that queues are equally weighted, except at ingress for mixed-speed LAGs without per-fp-inq-queuing enabled and at egress for all mixed-speed LAGs, in which cases the queues are weighted based on port speed. force-equal — Specifies that the queues are always equally weighted. offered-load — Specifies that the queues are weighted based on observed offered load. capped-offered-load — Specifies that the queues are weighted based on observed offered load capped. by PIR

hsmdb-pools

Syntax **hsmdb-pools mda** *mda-id* {**ingress** | **egress**} [**detail**]
hsmdb-pools port *port-id* **class** *class-id*
hsmdb-pools port *port-id* **queue-group** *queue-group-name* **instance** *instance-id*
hsmdb-pools sap *sap-id* [**port** *port-id*]
hsmdb-pools [**port** *port-id*] **subscriber** *sub-ident*
hsmdb-pools port *port-id* **network-queues**

Context show>qos

Description This command displays HSMDB pool information.

Parameters **detail** — Displays detailed information.
mda *mda-id* — Displays HSMDB pool information associated with the specified XMA or MDA.
Values slot/mda
ingress — Displays HSMDB ingress pool information associated with the specified XMA or MDA.
egress — Displays HSMDB egress pool information associated with the specified XMA or MDA.
port-id — Displays HSMDB pool information associated with the specified port.

class-id — Specifies the class-id.

sap-id — Specifies the SAP assigned to the service.

queue-group-name — Displays information about the specified queue group on the port.

instance-id — Specifies the identification of a specific instance of the queue group.

sub-ident — Specifies the subscriber identification.

Output

Sample Output

```
*A:Dut-A# show qos hsmdda-pools ingress mda 4/1
=====
Root Pools
=====
Pool ID          Size          Remaining
-----
1                117891        117889
2                117891        117891
3                117891        117885
4                117891        117891
5                117891        117891
6                117891        117891
7                117891        117891
8                117891        117891
=====
Class Pools
=====
Pool ID          Size          Remaining
-----
1                117891        117891
2                117891        117891
3                117891        117891
4                117891        117891
5                117891        117891
6                117891        117891
7                117891        117890
8                117891        117891
=====
Port Class Pools
=====
Port ID   Class ID   Size          Remaining
-----
1         1         117891        117884
1         2         117891        117891
1         3         117891        117891
1         4         117891        117891
1         5         117891        117891
1         6         117891        117891
1         7         117891        117891
1         8         117891        117889
=====
*A:Dut-A#

*A:Dut-A# show qos hsmdda-pools egress mda 4/1
```

```

=====
Root Pools
=====
Pool ID          Size          Remaining
-----
1                117891         117891
2                117891         117891
3                117891         117891
4                117891         117891
5                117891         117891
6                117891         117891
7                117891         117891
8                117891         117891
=====

Class Pools
=====
Pool ID          Size          Remaining
-----
1                117891         117891
2                117891         117884
3                117891         117891
4                117891         117884
5                117891         117891
6                117891         117891
7                117891         117891
8                117891         117891
=====

Port Class Pools
=====
Port ID  Class ID  Size          Remaining
-----
1        1        117891         117891
1        2        117891         117891
1        3        117891         117886
1        4        117891         117891
1        5        117891         117891
1        6        117891         117891
1        7        117891         117891
1        8        117891         117891
=====

*A:Dut-A#

*A:Dut-A# show qos hsmda-pools ingress mda 4/1 detail
Buffer Pools HSM DA 4/1

Port Allocation Factors
  Port 1 Act-BW: 10 Gbps Modifier: 100 Actual Factor: 10000

Root Pools Percentage of Total: 89.99 Actual Total Size 943128
  Pool 1      Allocation Percentage: 12.50 Size: 117891 Remaining: 117891

      Class-pool 1 Percentage of Root: 100.00 Size: 117891 Remaining: 117891
        Port-class-pool 1      Factor-Size: 117891 Remaining: 117891
        Port-class-pool 2      Factor-Size: 0 Remaining: 0
        Port-class-pool 3      Factor-Size: 0 Remaining: 0
        Port-class-pool 4      Factor-Size: 0 Remaining: 0
        Port-class-pool 5      Factor-Size: 0 Remaining: 0
        Port-class-pool 6      Factor-Size: 0 Remaining: 0

```

Port-class-pool	7	Factor-Size:	0	Remaining:	0
Port-class-pool	8	Factor-Size:	0	Remaining:	0
Port-class-pool	9	Factor-Size:	0	Remaining:	0
Port-class-pool	10	Factor-Size:	0	Remaining:	0
Pool 2	Allocation Percentage:	12.50	Size:	117891	Remaining:
					117884
Class-pool 2	Percentage of Root:	100.00	Size:	117891	Remaining:
					117891
Port-class-pool	1	Factor-Size:	117891	Remaining:	117891
Port-class-pool	2	Factor-Size:	0	Remaining:	0
Port-class-pool	3	Factor-Size:	0	Remaining:	0
Port-class-pool	4	Factor-Size:	0	Remaining:	0
Port-class-pool	5	Factor-Size:	0	Remaining:	0
Port-class-pool	6	Factor-Size:	0	Remaining:	0
Port-class-pool	7	Factor-Size:	0	Remaining:	0
Port-class-pool	8	Factor-Size:	0	Remaining:	0
Port-class-pool	9	Factor-Size:	0	Remaining:	0
Port-class-pool	10	Factor-Size:	0	Remaining:	0
Pool 3	Allocation Percentage:	12.50	Size:	117891	Remaining:
					117891
Class-pool 3	Percentage of Root:	100.00	Size:	117891	Remaining:
					117891
Port-class-pool	1	Factor-Size:	117891	Remaining:	117891
Port-class-pool	2	Factor-Size:	0	Remaining:	0
Port-class-pool	3	Factor-Size:	0	Remaining:	0
Port-class-pool	4	Factor-Size:	0	Remaining:	0
Port-class-pool	5	Factor-Size:	0	Remaining:	0
Port-class-pool	6	Factor-Size:	0	Remaining:	0
Port-class-pool	7	Factor-Size:	0	Remaining:	0
Port-class-pool	8	Factor-Size:	0	Remaining:	0
Port-class-pool	9	Factor-Size:	0	Remaining:	0
Port-class-pool	10	Factor-Size:	0	Remaining:	0
Pool 4	Allocation Percentage:	12.50	Size:	117891	Remaining:
					117891
Class-pool 4	Percentage of Root:	100.00	Size:	117891	Remaining:
					117891
Port-class-pool	1	Factor-Size:	117891	Remaining:	117891
Port-class-pool	2	Factor-Size:	0	Remaining:	0
Port-class-pool	3	Factor-Size:	0	Remaining:	0
Port-class-pool	4	Factor-Size:	0	Remaining:	0
Port-class-pool	5	Factor-Size:	0	Remaining:	0
Port-class-pool	6	Factor-Size:	0	Remaining:	0
Port-class-pool	7	Factor-Size:	0	Remaining:	0
Port-class-pool	8	Factor-Size:	0	Remaining:	0
Port-class-pool	9	Factor-Size:	0	Remaining:	0
Port-class-pool	10	Factor-Size:	0	Remaining:	0
Pool 5	Allocation Percentage:	12.50	Size:	117891	Remaining:
					117891
Class-pool 5	Percentage of Root:	100.00	Size:	117891	Remaining:
					117891
Port-class-pool	1	Factor-Size:	117891	Remaining:	117891
Port-class-pool	2	Factor-Size:	0	Remaining:	0
Port-class-pool	3	Factor-Size:	0	Remaining:	0
Port-class-pool	4	Factor-Size:	0	Remaining:	0
Port-class-pool	5	Factor-Size:	0	Remaining:	0
Port-class-pool	6	Factor-Size:	0	Remaining:	0
Port-class-pool	7	Factor-Size:	0	Remaining:	0
Port-class-pool	8	Factor-Size:	0	Remaining:	0
Port-class-pool	9	Factor-Size:	0	Remaining:	0
Port-class-pool	10	Factor-Size:	0	Remaining:	0
Pool 6	Allocation Percentage:	12.50	Size:	117891	Remaining:
					117891
Class-pool 6	Percentage of Root:	100.00	Size:	117891	Remaining:
					117889


```

Port-class-pool 1      Factor-Size: 117891  Remaining: 117891
Port-class-pool 2      Factor-Size: 0        Remaining: 0
Port-class-pool 3      Factor-Size: 0        Remaining: 0
Port-class-pool 4      Factor-Size: 0        Remaining: 0
Port-class-pool 5      Factor-Size: 0        Remaining: 0
Port-class-pool 6      Factor-Size: 0        Remaining: 0
Port-class-pool 7      Factor-Size: 0        Remaining: 0
Port-class-pool 8      Factor-Size: 0        Remaining: 0
Port-class-pool 9      Factor-Size: 0        Remaining: 0
Port-class-pool 10     Factor-Size: 0        Remaining: 0
Pool 7      Allocation Percentage: 12.50  Size: 117891  Remaining: 117891

Class-pool 7 Percentage of Root: 100.00  Size: 117891  Remaining: 117891
Port-class-pool 1      Factor-Size: 117891  Remaining: 117889
Port-class-pool 2      Factor-Size: 0        Remaining: 0
Port-class-pool 3      Factor-Size: 0        Remaining: 0
Port-class-pool 4      Factor-Size: 0        Remaining: 0
Port-class-pool 5      Factor-Size: 0        Remaining: 0
Port-class-pool 6      Factor-Size: 0        Remaining: 0
Port-class-pool 7      Factor-Size: 0        Remaining: 0
Port-class-pool 8      Factor-Size: 0        Remaining: 0
Port-class-pool 9      Factor-Size: 0        Remaining: 0
Port-class-pool 10     Factor-Size: 0        Remaining: 0
Pool 8      Allocation Percentage: 12.50  Size: 117891  Remaining: 117891

Class-pool 8 Percentage of Root: 100.00  Size: 117891  Remaining: 117884
Port-class-pool 1      Factor-Size: 117891  Remaining: 117891
Port-class-pool 2      Factor-Size: 0        Remaining: 0
Port-class-pool 3      Factor-Size: 0        Remaining: 0
Port-class-pool 4      Factor-Size: 0        Remaining: 0
Port-class-pool 5      Factor-Size: 0        Remaining: 0
Port-class-pool 6      Factor-Size: 0        Remaining: 0
Port-class-pool 7      Factor-Size: 0        Remaining: 0
Port-class-pool 8      Factor-Size: 0        Remaining: 0
Port-class-pool 9      Factor-Size: 0        Remaining: 0
Port-class-pool 10     Factor-Size: 0        Remaining: 0
*A:Dut-A#

```

hsmda-scheduler-hierarchy

Syntax **hsmda-scheduler-hierarchy** port *port-id* [*interval time*]
hsmda-scheduler-hierarchy port *port-id* {*shapers* | *shaper shaper-name*} [*interval time*]
[*detail*]
hsmda-scheduler-hierarchy mda *mda-id* [*interval time*]
hsmda-scheduler-hierarchy sap *sap-id* [*ingress* | *egress*] [*interval time*]
hsmda-scheduler-hierarchy subscriber *sub-id* [*ingress* | *egress*] [*interval time*]

Context show>qos

Description This command displays HSM DA scheduler hierarchy information.

Parameters **port** *port-id* — Displays HSMDA scheduler hierarchy information about the specified port.

Values

<i>port-id</i>	<i>slot/mda/port</i> [.channel]	
	<i>slot</i> [/ <i>mda</i> [/ <i>port</i>]]	
<i>eth-sat-id</i>	<i>esat-id/slot/port</i>	
	<i>esat</i>	keyword
	<i>id</i>	1 to 20
<i>pxc-id</i>	<i>pxc-id.sub-port</i>	
	<i>pxc</i>	keyword
	<i>id</i>	1 to 64
	<i>sub-port</i>	a, b
<i>aps-id</i>	<i>aps-group-id</i> [.channel]	
	<i>aps</i>	keyword
	<i>group-id</i>	1 to 64
<i>ccag-id</i>	<i>slot/mda/path-id</i> [<i>cc-type</i>]	
	<i>path-id</i>	a, b
	<i>cc-type</i>	.sap-net, .net-sap

shapers — Displays all shaper information.

shaper *shape-name* — Displays information about the specified shaper.

sap *sap-id* — Displays information about the specified SAP ID.

Values

<i>sap-id</i>	null	<i>port-id</i> <i>lag-id</i>
		<i>dot1q</i> <i>port-id</i> <i>lag-id</i> :qtag1
		<i>qinq</i> <i>port-id</i> <i>lag-id</i> :qtag1.qtag2
		<i>port-id</i> <i>slot/mda/port</i> [.channel]
		<i>esat-id/slot/port</i>
		<i>esat</i> keyword
		<i>id</i> 1 to 20
		<i>pxc-id.sub-port</i>
		<i>pxc</i> keyword
		<i>id</i> 1 to 64
		<i>sub-port</i> a, b
		<i>lag-id</i> <i>lag-id</i>
		<i>lag</i> keyword
		<i>id</i> 1 to 800

qtag1 *, 0 to 4094
qtag2 *, 0 to 4094

ingress — Displays ingress information about the SAP or subscriber.

egress — Displays egress information about the SAP or subscriber.

subscriber *sub-id* — Displays information about the specified subscriber.

Output

Sample Output

```
*A:Dut-A# show qos hsmda-scheduler-hierarchy port 4/1/1 shapers

HSMda Scheduler Policy egrSchedPol
  Port Bandwidth: 10 Gbps
  Max Rate      : 5000457120

Scheduler Priority 8  Rate: 1221455600
Scheduler Class 8   Rate: 1221455600

Scheduler Priority 7  Rate: 1221539520
Scheduler Class 7   Rate: 1221539520

Scheduler Priority 6  Rate: 1221623440
Scheduler Class 6   Rate: 1221623440

Scheduler Priority 5  Rate: 1221539520
Scheduler Class 5   Rate: 1221539520

Scheduler Priority 4  Rate: 113963360
Scheduler Class 4   Rate: 113963360

Scheduler Priority 3  Rate: 0
Scheduler Class 3   Rate: 0

Scheduler Priority 2  Rate: 0
Scheduler Class 2   Rate: 0

Scheduler Priority 1  Rate: 0
Scheduler Class 1   Rate: 0

Secondary Shaper secShaper  Rate 4999785760

*A:Dut-A#

*A:Dut-A# show qos hsmda-scheduler-hierarchy port 4/1/1 shaper secShaper

HSMda Scheduler Policy egrSchedPol
  Port Bandwidth: 10 Gbps
  Max Rate      : 4999785760

Scheduler Priority 8  Rate: 1221539520
```

```
Scheduler Class 8 Rate: 1221539520

Scheduler Priority 7 Rate: 1221623440
Scheduler Class 7 Rate: 1221623440

Scheduler Priority 6 Rate: 1221455600
Scheduler Class 6 Rate: 1221455600

Scheduler Priority 5 Rate: 1221455600
Scheduler Class 5 Rate: 1221455600

Scheduler Priority 4 Rate: 113963360
Scheduler Class 4 Rate: 113963360

Scheduler Priority 3 Rate: 0
Scheduler Class 3 Rate: 0

Scheduler Priority 2 Rate: 0
Scheduler Class 2 Rate: 0

Scheduler Priority 1 Rate: 0
Scheduler Class 1 Rate: 0

Secondary Shaper secShaper Rate 4999869680

*A:Dut-A#

*A:Dut-A# show qos hsmda-scheduler-hierarchy mda 4/1

HSM DA Scheduler Policy ingSchedPol
Max Rate : 9772064400

Scheduler Priority 8 Rate: 1221455600
Scheduler Class 8 Rate: 1221455600

Scheduler Priority 7 Rate: 1221455600
Scheduler Class 7 Rate: 1221455600

Scheduler Priority 6 Rate: 1221455600
Scheduler Class 6 Rate: 1221455600

Scheduler Priority 5 Rate: 1221539520
Scheduler Class 5 Rate: 1221539520

Scheduler Priority 4 Rate: 1221539520
Scheduler Class 4 Rate: 1221539520

Scheduler Priority 3 Rate: 1221539520
Scheduler Class 3 Rate: 1221539520

Scheduler Priority 2 Rate: 1221455600
Scheduler Class 2 Rate: 1221455600

Scheduler Priority 1 Rate: 1221539520
Scheduler Class 1 Rate: 1221539520

*A:Dut-A#
*A:Dut-A#
```

```
*A:Dut-A# show qos hsmda-scheduler-hierarchy sap 4/1/1:1 ingress
```

```
HSMDA Scheduler Policy ingSchedPol
Port Bandwidth: 10 Gbps
Max Rate       : 9772094400
```

```
Scheduler Priority 8 Rate: 1221539520
Scheduler Class 8 Rate: 1221539520
[4/1/1:1] Queue 8 Rate 14982240
```

```
Scheduler Priority 7 Rate: 1221539520
Scheduler Class 7 Rate: 1221539520
[4/1/1:1] Queue 7 Rate 14982240
```

```
Scheduler Priority 6 Rate: 1221455600
Scheduler Class 6 Rate: 1221455600
[4/1/1:1] Queue 6 Rate 14982240
```

```
Scheduler Priority 5 Rate: 1221539520
Scheduler Class 5 Rate: 1221539520
[4/1/1:1] Queue 5 Rate 14899920
```

```
Scheduler Priority 4 Rate: 1221455600
Scheduler Class 4 Rate: 1221455600
[4/1/1:1] Queue 4 Rate 14982240
```

```
Scheduler Priority 3 Rate: 1221455600
Scheduler Class 3 Rate: 1221455600
[4/1/1:1] Queue 3 Rate 14982240
```

```
Scheduler Priority 2 Rate: 1221539520
Scheduler Class 2 Rate: 1221539520
[4/1/1:1] Queue 2 Rate 14982240
```

```
Scheduler Priority 1 Rate: 1221539520
Scheduler Class 1 Rate: 1221539520
[4/1/1:1] Queue 1 Rate 14982240
```

```
*A:Dut-A#
```

```
*A:Dut-A# show qos hsmda-scheduler-hierarchy sap 4/1/1:1 egress
```

```
HSMDA Scheduler Policy egrSchedPol
Port Bandwidth: 10 Gbps
Max Rate       : 5000373200
```

```
Scheduler Priority 8 Rate: 1221623440
Scheduler Class 8 Rate: 1221623440
[4/1/1:1] Queue 8 Rate 14982240
```

```
Scheduler Priority 7 Rate: 1221455600
Scheduler Class 7 Rate: 1221455600
[4/1/1:1] Queue 7 Rate 14982240
```

```
Scheduler Priority 6 Rate: 1221539520
Scheduler Class 6 Rate: 1221539520
[4/1/1:1] Queue 6 Rate 14982240
```

```
Scheduler Priority 5 Rate: 1221455600
Scheduler Class 5 Rate: 1221455600
```

```
[4/1/1:1] Queue 5 Rate 14982240

Scheduler Priority 4 Rate: 113963360
Scheduler Class 4 Rate: 113963360
[4/1/1:1] Queue 4 Rate 1399440

Scheduler Priority 3 Rate: 0
Scheduler Class 3 Rate: 0
[4/1/1:1] Queue 3 Rate 0

Scheduler Priority 2 Rate: 0
Scheduler Class 2 Rate: 0
[4/1/1:1] Queue 2 Rate 0

Scheduler Priority 1 Rate: 0
Scheduler Class 1 Rate: 0
[4/1/1:1] Queue 1 Rate 0

*A:Dut-A#

*A:Dut-A# show qos hsmda-scheduler-hierarchy subscriber s1.1.1 egress
HSMda Scheduler Policy egrSchedPol
Port Bandwidth: 10 Gbps
Max Rate: 7972904400

Scheduler Priority 8 Rate: 1084307200
Scheduler Class 8 Rate: 1084307200
[4/2/1:41] Queue 8Rate 0

Scheduler Priority 7 Rate: 1085220800
Scheduler Class 7 Rate: 1085220800
[4/2/1:41] Queue 7Rate 112000

Scheduler Priority 6 Rate: 1018748800
Scheduler Class 6 Rate: 1018748800
[4/2/1:41] Queue 6Rate 0

Scheduler Priority 5 Rate: 1006536000
Scheduler Class 5 Rate: 1006536000
[4/2/1:41] Queue 5Rate 0

Scheduler Priority 4 Rate: 1002505600
Scheduler Class 4 Rate: 1002505600
[4/2/1:41] Queue 4Rate 0

Scheduler Priority 3 Rate: 951526400
Scheduler Class 3 Rate: 951526400
[4/2/1:41] Queue 3Rate 0

Scheduler Priority 2 Rate: 938782400
Scheduler Class 2 Rate: 938782400
[4/2/1:41] Queue 2Rate 57600

Scheduler Priority 1 Rate: 916656400
Scheduler Class 1 Rate: 916656400
[4/2/1:41] Queue 1Rate 51520

*A:Dut-A#
```

hsmda-scheduler-policy

Syntax	hsmda-scheduler-policy [<i>hsmda-scheduler-policy-name</i>] [associations] [detail]
Context	show>qos
Description	This command displays HSMDA scheduler policy information.
Parameters	<i>hsmda-scheduler-policy-name</i> — Displays information about the specified HSMDA scheduler policy. associations — Displays entities associated with the specified HSMDA scheduler policy. detail — Displays detailed information.
Output	

Sample Output

```
*A:Dut-A# show qos hsmda-scheduler-policy ingSchedPol detail
=====
QoS HSMDA Scheduler Policy
=====
Policy Name           : ingSchedPol
=====
Description           : Scheduler Policy Id ingSchedPol

Max Rate: max

Group 1 Rate: max
  No classes assigned to group

Group 2 Rate: max
  No classes assigned to group

No Group
  Class 1 - Rate: max
  Class 2 - Rate: max
  Class 3 - Rate: max
  Class 4 - Rate: max
  Class 5 - Rate: max
  Class 6 - Rate: max
  Class 7 - Rate: max
  Class 8 - Rate: max
-----
Associations
-----
- MDA Ingress : 4/1 override
=====
*A:Dut-A#

*A:Dut-A# show qos hsmda-scheduler-policy ingSchedPol association
=====
QoS HSMDA Scheduler Policy
=====
```

```

Policy Name          : ingSchedPol
=====
Description          : Scheduler Policy Id ingSchedPol

-----
Associations
-----
- MDA Ingress : 4/1 override
=====
*A:Dut-A#

```

hsmda-slope-policy

- Syntax** **hsmda-slope-policy** [*hsmda-slope-policy-name*] [**associations**] [**detail**]
- Context** clear>qos
- Description** This command displays HSMDB slope policy information.
- Parameters** *hsmda-slope-policy-name* — Displays information about the specified HSMDB slope policy.
- associations** — Displays entities associated with the specified HSMDB slope policy.
- detail** — Displays detailed information.

Output

Sample Output

```

*A:Dut-A# show qos hsmda-slope-policy slopePol detail
=====
Qos HSMDB Slope Policy
=====
Policy Name          : slopePol
=====
Description          : (Not Specified)
Provisioned Queue MBS : 168000
-----
High Slope Parameters
-----
Start Depth          : 100.00          Admin State : Enabled
Max Depth            : 100.00          Max Prob.   : 100.00
-----
Low Slope Parameters
-----
Start Depth          : 90.00          Admin State : Enabled
Max Depth            : 90.00          Max Prob.   : 100.00
-----
SAP-Ingress Associations
-----
Policy ID            Queues
-----
1000                  1 2 3 4 5 6 7 8

```

SAP-Egress Associations

Policy ID Queues

No Association Found.

SAP Override Associations

SAP ID Direction Queues

No Association Found.

Sub-Profile Override Associations

Sub-Profile Direction Queues

No Association Found.
=====

*A:Dut-A#

*A:Dut-A# show qos hsmda-slope-policy slopePol association

=====

Qos HSMda Slope Policy

=====

Policy Name : slopePol

=====

Description : (Not Specified)

SAP-Ingress Associations

Policy ID Queues

1000 1 2 3 4 5 6 7 8

SAP-Egress Associations

Policy ID Queues

No Association Found.

SAP Override Associations

SAP ID Direction Queues

No Association Found.

Sub-Profile Override Associations

Sub-Profile Direction Queues

No Association Found.
=====

*A:Dut-A#

15 High Scale QoS IOM

15.1 Introduction

This section describes the capabilities of the High Scale QoS (HSQ) IOM. The HSQ IOM is an FP3-based IOM that has a multicore CPU and accepts up to two MDA-e cards. The HSQ IOM supports an enhanced egress QoS traffic manager architecture to provide scalable network, service, and subscriber QoS. At ingress, the HSQ IOM supports regular FP3 QoS with a high ingress policer scaling.

The HSQ IOM (iom4-e-hs) supports six scheduling classes across multiple hierarchical levels of hardware egress shaping with very stringent egress burst control. The scheduling allows a mix of strict priority and weighted round-robin (WRR) queues. A flexible buffer pool structure permits both buffer isolation and buffer oversubscription for the queue buffer allocation.

The HSQ IOM supports 768k queues, which are grouped into 96k queue groups; each comprises eight queues (referred to as HSQ queue groups). HSQ queue groups are used for SAP and subscriber egress queues, network egress queues, and both access and network egress queue group instance queues.

The components of the HSQ IOM operation are follows:

- [HSQ Shaping](#)
- [HSQ Scheduling](#)
- [HSQ Buffer Management](#)
- [HSQ LAGs](#)

15.1.1 HSQ Shaping

The HSQ egress shaping uses the following objects:


- HSQ queue groups — An HSQ queue group comprises eight egress queues with two WRR groups. One HS queue group is allocated to each of the following:
 - an egress SAP
 - an egress network port
 - an egress access queue group instance
 - an egress network queue group instance

- a subscriber egress (a single SLA profile instance in single HS SLA mode)
- primary shapers — A primary shaper is allocated for each secondary shaper because it is required in the hierarchy. It does not perform any QoS control except when allocated for subscriber egress configured with multiple SLA profile instances in extended HS SLA mode.
- secondary shapers — Secondary shapers provide an abstraction to be used for QoS control of traffic to a downstream device, such as an access node. Shaping can be performed on all traffic or on each scheduling class within the secondary shaper.
- ports — The traffic forwarded to each port can be shaped. In addition, traffic in each scheduling class within a port can be shaped individually or within a single WRR group.

Six scheduling classes are supported across all the preceding objects.

The egress QoS scheduling hierarchy is shown in [Figure 50](#).

Figure 50 Egress HSQ IOM Scheduling Hierarchy



sw0635

The available egress shaping encompasses:

- Per-queue or per-WRR group of queues

- per-HSQ queue group aggregate
- per-primary shaper aggregate
- per-secondary shaper aggregate
- per-secondary shaper per scheduling class
- per-port aggregate
- per-port per scheduling class

15.1.2 HSQ Scheduling

HSQ scheduling allows a combination of strict priority and WRR. There are six scheduling classes which are implemented from the HSQ queue group queues through the primary shaper, secondary shaper, and port. The scheduling classes are serviced in a strict priority order (scheduling class 6 having the highest priority and scheduling class 1 having the lowest priority), with WRR groups at the HSQ queue group and port levels, and a dynamic weight at the primary and secondary shaper levels.

Packet forwarding is achieved using service lists. The objects at each level are on a service list at the same level if they are in a state ready to send packets. The objects are off the service list if they have exceeded their configured PIR together with its related burst. When a port has a scheduling opportunity, it selects the secondary shaper to be serviced next, which selects the primary shaper to be serviced next, which selects the HSQ queue group to be serviced next, which selects a queue to be serviced next, resulting in a packet from that queue being forwarded.

At the HSQ queue group level, queues can be attached to one of two WRR groups. Each group is scheduled at a single scheduling class, with packets being taken from the constituent queues based on a configured queue weighting.

Weighting is also supported between queues and WRR groups in different HSQ queue groups per-primary shaper scheduling class. This allows the capacity available at the primary shaper scheduling class to be shared in a WRR manner between the HSQ queue group queues and WRR groups attached to that scheduling class.

There is a single WRR group at the port level that allows multiple scheduling classes to be collapsed to a single class per port, with each class in the group being assigned a weight.


The dynamic weights at the primary and secondary shapers are managed by the system, based on the number of pending packets for each of the shapers, not on the number of attached objects in each. The more pending packets a shaper has, the higher the weight it is assigned. The goal of the scheduling is to ensure a balanced distribution of capacity between each of the primary shapers and each of the secondary shapers. For example, this allows a secondary shaper with 10,000 active HSQ queue groups to receive proportionately more scheduling opportunities than another secondary shaper with only 100 active HSQ queue groups.

15.1.3 HSQ Buffer Management

The HSQ supports a flexible buffer management configuration that allows both buffer isolation and buffer oversubscription for the queue buffer allocation. There are four levels to the buffer hierarchy with the total buffer allocation being divided into a system-reserved portion and a user-provisioned portion, as shown in [Figure 51](#).

- root pools
- mid pools
- port class pools
- queue group queues

Figure 51 HSQ Buffer Pool Hierarchy



sw0636

15.1.4 HSQ LAGs

LAGs are supported on HSQ ports. The LAG port-type must be set to **hs** to add an HSQ port to a LAG, at which point only HSQ ports can be added to that LAG. When an HSQ queue group is created on a LAG, an HSQ queue group is allocated on each LAG port.

15.2 HSQ IOM-Specific Configurations

The HSQ IOM functionality is configured using the following set of HSQ-specific policies:

- HS attachment policy — Configures how HSQ queue group queues and WRR groups within SAP egress QoS policies, network queue policies, and egress queue group templates are attached to a scheduling class managed by the port scheduler
- HS pool policy — Configures the buffer allocation and hierarchy for the root and mid pools on the HSQ IOM
- HS port pool policy — Configures the buffer allocation and hierarchy for the port standard and alternative class pools
- HS scheduler policy — Configures the scheduling behavior of the HSQ scheduler classes at a port

The HSQ queue groups are configured using SAP egress QoS policies, network queue policies, and egress queue group templates, in which there are HSQ-specific commands corresponding to the HSQ functionality.

The HS secondary shapers, which provide QoS control for traffic to a downstream device, are configured directly in the port context on which the QoS control is required.

Overrides are available under ports and SAPs for the HSQ-specific configuration.

See to the High Scale QoS IOM Command Reference section for details about command syntax and descriptions.

15.3 Non-Applicable HSQ IOM Functions

As the HSQ IOM traffic manager replaces several egress FP QoS functions, the following configurations are not applicable to an HSQ IOM.

- QoS related
 - advanced configuration policies
 - egress access and network MDA and port pools
 - egress **agg-rate queue-frame-based-accounting**
 - egress H-QoS virtual scheduling, including Vports
 - FP egress queue
 - drop tails and CBS
 - WRED per queue (an HSQ equivalent is available)
 - CIR
 - **highplus** slope in a slope policy
 - burst limit (an HSQ equivalent is available)
 - dynamic MBS
 - network queue MBS (replaced by the **hs-mbs** command)
 - HS-MDAv2
 - multi-service sites
 - PBB egress B-SAP per ISID shaping
- subscriber related
 - Access Node Control Protocol Management (ANCP)
 - PCC rule-based subscriber services (Gx or RADIUS)
 - Web Authentication Protocol (WPP)
- MPLS related
 - Generalized Multiprotocol Label Switching (GMPLS) UNI
- Service related
 - G.8031 protected Ethernet tunnels
- System related
 - Ethernet satellite host ports
 - port cross-connects (PXC)
 - Sflow

15.4 High Scale QoS IOM Command Reference

15.4.1 Command Hierarchies

- [HS Attachment Policy Commands](#)
- [HS Pool Policy Commands](#)
- [HS Port Pool Policy Commands](#)
- [HS Scheduler Policy Commands](#)
- [Network Queue QoS Policy HS Commands](#)
- [SAP Egress Policy QoS HS Commands](#)
- [Queue Group Template HS Commands](#)
- [Services Configuration HS Commands](#)
- [Forwarding Plane HS Commands](#)
- [Port Configuration HS Commands](#)
- [LAG HS Commands](#)
- [Operational Commands](#)
- [Show Commands](#)
- [Clear Commands](#)

15.4.1.1 HS Attachment Policy Commands

```
config
— qos
  — hs-attachment-policy policy-name [create]
  — no hs-attachment-policy policy-name
    — description description-string
    — no description
    — low-burst-max-class class
    — no low-burst-max-class
    — queue queue-id sched-class class-id
    — queue queue-id unattached
    — queue queue-id wrr-group wrr-group-id
    — no queue queue-id
    — wrr-group group-id sched-class class-id
    — wrr-group group-id unattached
    — no wrr-group group-id
```

15.4.1.2 HS Pool Policy Commands

```

config
  — qos
    — hs-pool-policy policy-name [create]
    — no hs-pool-policy policy-name
      — description description-string
      — no description
      — mid-tier
        — [no] mid-pool mid-pool-id
          — allocation-percent percent-of-parent-pool
          — no allocation-percent
          — parent-root-pool root-pool-id
          — no parent-root-pool
          — port-bw-oversub-factor oversubscription-factor
          — no port-bw-oversub-factor
          — slope-policy policy-name
          — no slope-policy
        — root-tier
          — [no] root-pool root-pool-id
            — allocation-weight pool-weight
            — no allocation-weight
            — slope-policy policy-name
            — no slope-policy
      — system-reserve percent-of-buffers
      — no system-reserve

```

15.4.1.3 HS Port Pool Policy Commands

```

config
  — qos
    — hs-port-pool-policy policy-name [create]
    — no hs-port-pool-policy policy-name
      — alt-port-class-pools
        — [no] class-pool alt-class-pool-id
          — allocation explicit-percent percent-of-parent-pool
          — allocation port-bw-weight pool-weight
          — no allocation
          — parent-mid-pool mid-pool-id
          — no parent-mid-pool
          — slope-policy policy-name
          — no slope-policy
        — description description-string
        — no description
      — std-port-class-pools
        — [no] class-pool std-class-pool-id
          — allocation explicit-percent percent-of-parent-pool
          — allocation port-bw-weight pool-weight
          — no allocation
          — parent-mid-pool mid-pool-id

```

- no **parent-mid-pool**
- **slope-policy** *policy-name*
- no **slope-policy**

15.4.1.4 HS Scheduler Policy Commands

```
config
  — qos
    — hs-scheduler-policy policy-name [create]
    — no hs-scheduler-policy policy-name
      — description description-string
      — no description
      — group group-id rate rate
      — no group group-id
      — max-rate rate
      — no max-rate
      — scheduling-class class-id group group-id [weight weight-in-group]
      — scheduling-class class-id rate rate
      — no scheduling-class class-id
```

15.4.1.5 Network Queue QoS Policy HS Commands

```
config
  — qos
    — network-queue policy-name [create]
    — no network-queue policy-name
      — hs-attachment-policy policy-name
      — no hs-attachment-policy
      — [no] hs-wrr-group group-id
        — adaptation-rule [pir adaptation-rule]
        — no adaptation-rule
        — hs-class-weight weight
        — no hs-class-weight
        — rate percent
        — no rate
      — queue queue-id [multipoint] [queue-type] [create]
      — no queue queue-id
        — [no] hs-alt-port-class-pool
        — hs-class-weight weight
        — no hs-class-weight
        — hs-mbs percent-of-queue-rate
        — no hs-mbs
        — hs-wred-queue [policy slope-policy-name]
        — no hs-wred-queue
        — hs-wrr-weight weight
        — no hs-wrr-weight
```

15.4.1.6 SAP Egress Policy QoS HS Commands

```

config
  — qos
    — sap-egress {policy-id | policy-name} [create] [name name] [create]
    — no sap-egress {policy-id | policy-name}
      — hs-attachment-policy policy-name
      — no hs-attachment-policy
      — hs-wrr-group group-id
      — no hs-wrr-group
        — adaptation-rule [pir adaptation-rule]
        — no adaptation-rule
        — hs-class-weight weight
        — no hs-class-weight
        — percent-rate percent
        — no percent-rate
        — rate rate
        — no rate
      — queue queue-id [queue-type] [create]
      — no queue queue-id
        — [no] hs-alt-port-class-pool
        — hs-class-weight weight
        — no hs-class-weight
        — hs-wred-queue [policy slope-policy-name]
        — no hs-wred-queue
        — hs-wrr-weight weight
        — no hs-wrr-weight

```

15.4.1.7 Queue Group Template HS Commands

```

config
  — qos
    — queue-group-templates
      — egress
        — queue-group queue-group-name [create]
        — queue-group queue-group-name
          — hs-attachment-policy policy-name
          — no hs-attachment-policy
          — [no] hs-wrr-group group-id
            — adaptation-rule [pir adaptation-rule]
            — no adaptation-rule
            — hs-class-weight weight
            — no hs-class-weight
            — percent-rate percent
            — no percent-rate
            — rate rate
            — no rate
          — no queue queue-id
            — [no] hs-alt-port-class-pool
            — hs-class-weight weight

```

- **no** **hs-class-weight**
- **hs-wred-queue** [policy slope-policy-name]
- **no** **hs-wred-queue**
- **hs-wrr-weight** weight
- **no** **hs-wrr-weight**

15.4.1.8 Services Configuration HS Commands

The following services configuration commands are specific to HSQ IOMs. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN, 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more details.

- [Epipe Commands](#)
- [IES Commands](#)
- [Ipipe Commands](#)
- [VPLS Commands](#)
- [VPRN Commands](#)

15.4.1.8.1 Epipe Commands

```
config>service
  — epipe service-id service-id [customer customer-id] [create] [vpn vpn-id] [vc-switching] [test]
    [name name]
  — no epipe service-id
    — sap sap-id [create] [no-endpoint]
    — sap sap-id [create] endpoint endpoint-name
    — no sap sap-id
    — egress
      — [no] queue-override
        — hs-secondary-shaper policy-name
        — no hs-secondary-shaper
        — hs-wrr-group group-id [create]
        — no hs-wrr-group
          — class-weight weight
          — no class-weight
          — percent-rate percent
          — no percent-rate
          — rate rate
          — no rate
        — queue queue-id [create]no queue queue-id
          — hs-class-weight weight
          — no hs-class-weight
          — hs-wred-queue policy slope-policy-name
          — no hs-wred-queue
```

- **hs-wrr-weight** *weight*
- **no hs-wrr-weight**

15.4.1.8.2 IES Commands

```

config>service
— ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
— no ies service-id
  — [no] interface ip-int-name
    — [no] sap sap-id
      — egress
        — [no] queue-override
          — hs-secondary-shaper policy-name
          — no hs-secondary-shaper
          — hs-wrr-group group-id [create]
          — no hs-wrr-group group-id
            — class-weight weight
            — no class-weight
            — percent-rate percent
            — no percent-rate
            — rate rate
            — no rate
          — queue queue-id [create]
          — no queue queue-id
            — hs-class-weight weight
            — no hs-class-weight
            — hs-wred-queue policy slope-policy-name
            — no hs-wred-queue
            — hs-wrr-weight weight
            — no hs-wrr-weight

```

15.4.1.8.3 Ipipe Commands

```

config>service
— ipipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-switching] [name name]
— no ipipe service-id
  — sap sap-id [create] [no-endpoint]
  — sap sap-id [create] endpoint endpoint-name]
  — no sap sap-id
    — egress
      — [no] queue-override
        — hs-secondary-shaper policy-name
        — no hs-secondary-shaper
        — hs-wrr-group group-id [create]
        — no hs-wrr-group group-id
          — class-weight weight
          — no class-weight
          — percent-rate percent
          — no percent-rate

```


- **rate** *rate*
- **no rate**
- **queue** *queue-id* [**create**]**no queue** *queue-id*
 - **hs-class-weight** *weight*
 - **no hs-class-weight**
 - **hs-wred-queue** *policy slope-policy-name*
 - **no hs-wred-queue**
 - **hs-wrr-weight** *weight*
 - **no hs-wrr-weight**

15.4.1.8.4 VPLS Commands

```
config>service
— vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls] [etree] [name
name] [create]
— no vpls service-id
  — sap sap-id [create] [no-endpoint]
  — sap sap-id [create] endpoint endpoint-name
  — no sap sap-id
    — egress
      — [no] queue-override
        — hs-secondary-shaper policy-name
        — no hs-secondary-shaper
        — hs-wrr-group group-id [create]
        — no hs-wrr-group group-id
          — class-weight weight
          — no class-weight
          — percent-rate percent
          — no percent-rate
          — rate rate
          — no rate
        — queue queue-id [create]no queue queue-id
          — hs-class-weight weight
          — no hs-class-weight
          — hs-wred-queue policy slope-policy-name
          — no hs-wred-queue
          — hs-wrr-weight weight
          — no hs-wrr-weight
```

15.4.1.8.5 VPRN Commands

```
config>service
— vpn service-id [name name] [customer customer-id] [create]
— no vpn service-id
  — [no] interface ip-int-name
    — [no] sap sap-id
      — egress
        — [no] queue-override
          — hs-secondary-shaper policy-name
```

-
- no hs-secondary-shaper
 - hs-wrr-group *group-id* [create]
 - no hs-wrr-group *group-id*
 - class-weight *weight*
 - no class-weight
 - percent-rate *percent*
 - no percent-rate
 - rate *rate*
 - no rate
 - queue *queue-id* [create]
 - no queue *queue-id*
 - hs-class-weight *weight*
 - no hs-class-weight
 - hs-wred-queue policy *slope-policy-name*
 - no hs-wred-queue
 - hs-wrr-weight *weight*
 - no hs-wrr-weight

15.4.1.9 Forwarding Plane HS Commands

The following forwarding plane configuration commands are specific to HSQ IOMs. For more information, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide*.

```
config>card>fp>egress
  — hs-fixed-high-thresh-delta size-in-bytes
  — no hs-fixed-high-thresh-delta
  — hs-pool-policy name
  — no hs-pool-policy
```

15.4.1.10 Port Configuration HS Commands

The following port configuration commands are specific to HSQ IOMs. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide* for more details.

```
config>port>ethernet
  — access
    — egress
      — [no] queue-group queue-group-name [create] [instance instance-id]
        — [no] hs-turbo
    — egress
      — hs-port-pool-policy policy-name
      — no hs-port-pool-policy
      — hs-scheduler-overrides [create]
      — no hs-scheduler-overrides
        — group group-id rate rate
        — no group group-id
        — max-rate rate
        — no max-rate
        — scheduling-class class rate rate
        — scheduling-class class weight weight-in-group
        — no scheduling-class class
      — hs-scheduler-policy policy-name
      — no hs-scheduler-policy
      — hs-secondary-shaper secondary-shaper-name [create]
      — no hs-secondary-shaper secondary-shaper-name
        — [no] aggregate
          — low-burst-max-class class
          — no low-burst-max-class
          — rate rate
          — no rate
        — [no] class class-number
          — rate rate
          — no rate
        — description description-string
        — no description
  — network
    — egress
      — [no] queue-group queue-group-name [create] [instance instance-id]
        — [no] hs-turbo
```

15.4.1.11 LAG HS Commands

The following LAG configuration command has a parameter that is specific to HSQ IOMs. For further information, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide*.

```
config
```

- lag
 - port-type *lag-port-type*
 - no port-type

15.4.1.12 Operational Commands

- config
- qos
 - **copy** **hs-attachment-policy** *src-name dst-name* [overwrite]
 - **copy** **hs-pool-policy** *src-name dst-name* [overwrite]
 - **copy** **hs-port-pool-policy** *src-name dst-name* [overwrite]
 - **copy** **hs-scheduler-policy** *src-name dst-name* [overwrite]

15.4.1.13 Show Commands

- show
- **hs-pools** *card-slot-number fp forwarding-plane egress*
 - **hs-pools** **port** *port-id egress*
 - **hs-pools** **port** *port-id egress network-queues*
 - **hs-pools** **port** *port-id egress queue-group queue-group-name* [instance *instance-id*]
 - **hs-pools** **port** *port-id egress sap sap-id*
 - **hs-pools** **port** *port-id egress subscriber sub-ident-string*
 - **port** *port-id hs-secondary-shaper*
 - **port** *port-id hs-secondary-shaper secondary-shaper-name* [statistics | associations]
 - qos
 - **hs-attachment-policy** [*policy-name*] [association | detail]
 - **hs-pool-policy** [*policy-name*] [association | detail]
 - **hs-port-pool-policy** [*policy-name*] [association | detail]
 - **hs-scheduler-hierarchy** **port** *port-id* [hs-secondary-shaper *shaper-name*] [interval *time-in-seconds*]
 - **hs-scheduler-hierarchy** **port** *port-id* [interval *time-in-seconds*] **hs-secondary-shapers**
 - **hs-scheduler-hierarchy** **port** *port-id* [interval *time-in-seconds*] **queue-group** *queue-group-name* **instance** *instance-id* {access | network}
 - **hs-scheduler-hierarchy** **sap** *sap-id* **egress** [interval *time-in-seconds*]
 - **hs-scheduler-hierarchy** **subscriber** *sub-ident* **egress** [interval *time-in-seconds*]
 - **hs-scheduler-policy** [*policy-name*] [association | detail]

15.4.1.14 Clear Commands

- clear
- **card** *slot-number fp fp-number* **hs-pool** **high-water-mark**
 - **card** *slot-number fp fp-number* **hs-pool** **high-water-mark** **mid-pool** [*value*]
 - **card** *slot-number fp fp-number* **hs-pool** **high-water-mark** **root-pool** [*value*]
 - **card** *slot-number fp fp-number* **hs-pool** **high-water-mark** **system**

-
- **port** *port-id* **hs-secondary-shaper** *shaper-name* **statistics**
 - **port** *port-id* **hs-pool high-water-mark** {[*standard value*] | [*alternate value*]}

15.4.2 Command Descriptions

15.4.2.1 Configuration Commands

15.4.2.1.1 Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>hs-attachment-policy config>qos>hs-pool-policy config>qos>hs-port-pool-policy config>qos>hs-scheduler-policy
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	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

15.4.2.2 HS Attachment Policy Commands

hs-attachment-policy

Syntax	hs-attachment-policy <i>policy-name</i> [create] no hs-attachment-policy <i>policy-name</i>
Context	config>qos

Description	This command specifies how the queues within an HSQ queue group associated with the SAP egress policy instance, egress queue group instance, or egress network queue policy instance attaches to the HSQ scheduling classes managed by the port scheduler. On the HSQ IOM, eight queues are allocated per egress SAP or subscriber SLA profile instance (SPI), or per egress (access or network) queue group instance, or per egress network port, numbered 1 through 8. The port scheduler maintains six scheduling classes numbered from 1 through 6 (6 being the highest relative priority and 1 being the lowest). The set of eight queues may also be placed into one of two local Weighted Round Robin (WRR) groups which collapse the member queues into a single scheduling class while providing a weighted fair distribution of scheduling opportunities per member. The attachment policy contains the attachment commands that map the queue IDs and WRR groups to the scheduling classes. The attachment policy also defines the mapping of the scheduling classes to the queue's aggregate shapers low and high burst limit thresholds. The no form of the command deletes the HS attachment policy from the system, which is only possible if the policy is not being referenced.
Parameters	<i>policy-name</i> — Specifies an existing attachment policy, up to 32 characters. Each HS attachment policy must be uniquely named within the system. create — This keyword is required when first creating the configuration context. After the context is created, it is possible to navigate into the context without the create keyword.

low-burst-max-class

Syntax	low-burst-max-class <i>class</i> no low-burst-max-class
Context	config>qos>hs-attachment-policy
Description	This command specifies which scheduling classes map to the low burst-limit threshold of the queue-level aggregate shaper. Each egress SAP or subscriber SLA profile instance (SPI), per port set of network interface queues and egress queue group template instance has an aggregate shaper that manages the maximum burst limit over a specified shaping rate. Each aggregate shaper supports two thresholds. As the scheduling rate for the set of queues increases, eventually the aggregate rate exceeds the rate limit and the aggregate burst limit starts to be consumed. If this continues, the low burst limit threshold is exceeded and the queues mapped to the scheduling classes associated with low threshold are removed from the scheduler. If the remaining aggregate rate (from the higher scheduling classes) continues to exceed the shaping rate, then the burst limit continues to be consumed and eventually the high burst limit threshold is exceeded. This causes the queues for all scheduling classes to be removed from the scheduler.

The second (high) threshold exists to allow the higher priority classes to continue to forward, thereby mitigating the effects of low priority bursts beyond the aggregate shaping rate. Typically, the higher scheduling class queues are either individually rate-limited so their aggregate allowed throughput is less than the aggregate rate or the expected aggregate unshaped traffic from the individual higher scheduling classes does not exceed the aggregate shaping rate. Determining the value of **low-burst-max-class** involves anticipating the proper dividing line between the low and high scheduling classes by evaluating the forwarding behavior and SLA enforcement of each class.

By default, all scheduling classes are mapped to the low burst limit threshold. When mapping scheduling classes to the high burst limit threshold, an adequate value for the **card>fp>egress>hs-fixed-high-thresh-delta** must be specified. This is due to the fact that the queues associated with the lower classes may burst over the lower threshold in normal operation due to the scheduler forwarding whole packets. Set the **hs-fixed-highthresh-delta** value to at least two times the maximum frame size to prevent lower threshold class forwarding from also affecting the higher threshold classes when forwarding larger packet sizes. An insufficient high threshold delta defeats the intended purpose of mapping classes to the higher threshold.

The system utilizes the lowest value attainable for each low threshold aggregate burst limit without causing shaper underrun conditions. The high burst limit threshold is determined by adding the **hs-fixed-high-thresh-delta** value configured in the **config>card>fp>egress** CLI context to the aggregate's low burst limit threshold value.

The **low-burst-max-class** value can be changed at any time in the HS attachment policy. Modifying the setting causes all queue aggregate shapers to reconfigure the scheduling class mappings to the low and high burst limit thresholds to reflect the new value for scheduling class ID.

Scheduling Classes — As described in the **queue** and **wrr-group** attachment commands, each queue is either directly or indirectly (through a WRR group) mapped to a scheduling class. Each scheduling class has an inherent priority at the port scheduler. The inherent descending priority is as follows:

- Scheduling Class 6 (Highest)
- Scheduling Class 5
- Scheduling Class 4
- Scheduling Class 3
- Scheduling Class 2
- Scheduling Class 1 (lowest)

Placing scheduling classes into the port level WRR group causes those classes to compete for scheduling opportunities based on their associated weights instead of inherent priority. If higher weights are given to higher scheduling class IDs, then the relative proportional scheduling priority may continue to exhibit the priority level indicated by the class ID.

Setting Low and High Burst Limit Threshold Association — [Table 119](#) demonstrates the effect of the **low-burst-max-class** command parameters on scheduling class mappings to the low and high burst limit thresholds.

Table 119 Low and High Burst Limit Threshold Association

low-burst-max-class <i>sched-class</i>	Scheduling Classes on Low Threshold	Scheduling Classes on High Threshold
1	1	2, 3, 4, 5, and 6
2	1 and 2	3, 4, 5, and 6
3	1, 2, and 3	4, 5, and 6
4	1, 2, 3, and 4	5 and 6
5	1, 2, 3, 4, and 5	6
6	1, 2, 3, 4, 5, and 6	None

The **no** form of the command restores the queue aggregate low burst limit threshold maximum scheduling class mapping to the default value. This causes all sets of queues associated with the specified **hs-attachment-policy** *policy-name* to have all scheduling classes mapped to the low burst limit threshold.

Default low-burst-max-class 6

Parameters *class* — Specifies the low burst maximum class. This parameter is required when executing the **low-burst max-class** command. The parameter reflects the highest scheduling class that is associated with the low burst limit threshold associated with the queue aggregate shaper. Scheduling classes higher than the scheduling class ID are associated with the high burst limit threshold.

Values 1 to 6

queue

Syntax **queue** *queue-id* **sched-class** *class-id*
queue *queue-id* **unattached**
queue *queue-id* **wrr-group** *wrr-group-id*
no queue *queue-id*

Context config>qos>hs-attachment-policy

Description This command defines how the specified queue-id is attached to the scheduler. A queue may have one of four attachment states:

- directly attached to a scheduler class
- attached to a WRR group which is attached to a scheduler class
- attached to a WRR group which is unattached
- unattached

The following items are rules for attachment:

- Only one queue or WRR group can be attached to a given scheduling class. If another queue or a WRR group is currently attached to the specified scheduling class, the queue's attachment command fails and the current attachment for the *queue-id* is unchanged.
- Queues must be attached to scheduler classes (directly or indirectly through a WRR group) in an ascending order.
For example, if queue 3 is attached to scheduler class 2, queues 1 and 2 cannot be attached to scheduling classes 3 through 6 (or attached to a WRR group that is attached to scheduling classes 3 through 6).
- Up to six queues can be placed in any WRR group. These queues can all be in one WRR group or be spread between both WRR groups. Attempting to execute the **wrr-group** keyword on a seventh queue within the attachment policy fails.
- Queues must attach to WRR groups in contiguous order.
- WRR group 1 must have lower queue ID members than WRR group 2. Attempting to attach a queue ID to group 1 that is higher than any queue ID currently attached to group 2 fails with no change to the current attachment state for that queue ID. Attempting to attach a queue ID to group 2 that is lower than any queue ID currently attached to group 1 fails with no change to the current attachment state for that queue ID.
- WRR group 1 must have at least one attached member queue before queues can be attached to WRR group 2.
- Queues and WRR groups can be unattached at any time.
 - Unattached queues, or queues attached to an unattached WRR group, discard all received packets. Normal discard accounting is maintained.
 - The unattached state for an in-service queue or WRR group is intended for **hs-attachment-policy** editing purposes and is expected to be an intermediate state in this case.

When an **hs-attachment-policy** is initially created, all queues and both WRR groups default to the unattached state. Each queue and WRR group attachment state must be explicitly configured.

The **no** form of the command reverts to the default unattached attachment state for queue ID. The command fails if the specified *queue-id* is currently in a WRR group and removing the queue from that group causes the queue IDs for that group to become discontinuous.

Default	queue <i>queue-id</i> unattached
Parameters	<i>queue-id</i> — Specifies the queue identifier for the HS attachment policy queue. This parameter is required when executing the command. Values 1 to 8 sched-class — Specifies a direct attachment between the <i>queue-id</i> and one of the six scheduling classes. A value of 1 through 6 must accompany the sched-class keyword representing the queue's attached scheduling class. The sched-class, group, and unattached keywords are mutually exclusive. One of the keywords must be specified when the queue attachment command is executed.

class-id — Specifies the scheduling class that is associated with this *queue-id*.

Values 1 to 6

wrr-group — Specifies the inclusion of the *queue-id* in either WRR group 1 or WRR group 2. A value of 1 or 2 must accompany the **wrr-group** keyword representing the queue's attached WRR group number. Queues placed in a WRR group must have contiguous queue IDs. The **sched-class**, **group**, and **unattached** keywords are mutually exclusive. One of the keywords must be specified when the **queue** attachment command is executed.

wrr-group-id — Specifies the WRR group to which the *queue-id* is to be attached.

Values 1, 2

unattached — Indicates that the *queue-id* is not attached to any scheduling class or WRR group. An unattached queue does not forward any packets. The **sched-class**, **group**, and **unattached** keywords are mutually exclusive. One of the keywords must be specified when the **queue** attachment command is executed.

wrr-group

Syntax	wrr-group <i>group-id</i> sched-class <i>class-id</i> wrr-group <i>group-id</i> unattached no wrr-group <i>group-id</i>
Context	config>qos>hs-attachment-policy
Description	This command defines how the specified group ID is attached to the scheduler. A WRR group may have one of two attachment states: • directly attached to a scheduler class • unattached

A WRR group provides a weighted scheduling context for its member queues, collapsing the queues into a single scheduling class.

The following WRR membership restrictions apply:

- Two groups are supported: WRR-group 1 and WRR-group 2.
- Up to six of the eight queues can become members of WRR groups (the sum of the membership of both groups cannot exceed six queues). All six queues can be placed in a single WRR group or they can be spread between the two groups.
- All queues within a group must have contiguous queue IDs.
- When both groups are configured to have queue members, the queue IDs in group 1 must be lower than the queue IDs in group 2.

The **queue** *queue-id* attachment command is used to define WRR group membership.

The **no** form of the command reverts to the default unattached attachment state for the group ID.

Default	wrr-group group-id unattached
Parameters	<i>group-id</i> — Specifies the WRR group identifier. The <i>group-id</i> parameter is required when executing the wrr-group attachment command. Values 1, 2 sched-class — Specifies a direct attachment between the WRR group and one of the six scheduling classes. The sched-class and unattached keywords are mutually exclusive. One of the keywords must be specified when the wrr-group attachment command is executed. <i>class-id</i> — Specifies the scheduling class associated with this WRR group. The sched-class keyword specifies the attachment between the group ID and one of the six scheduling classes. A value of 1 through 6 must accompany the sched-class keyword representing the WRR group's attached scheduling class. Only one queue or WRR group can be attached to a given scheduling class. If another queue or a WRR group is currently attached to the specified scheduling class, the wrr-group attachment command fails and the current attachment for the <i>group-id</i> is unchanged. Values 1 to 6 unattached — Indicates that the group ID is not attached to any scheduling class or WRR group. Queues that are members of the unattached WRR group do not forward any packets. The sched-class and unattached keywords are mutually exclusive. One of the keywords must be specified when the wrr-group attachment command is executed.

15.4.2.3 HS Pool Policy Commands

hs-pool-policy

Syntax	hs-pool-policy <i>policy-name</i> [create] no hs-pool-policy <i>policy-name</i>
Context	config>qos
Description	This command enables the context to create HS pool policy parameters. The policy can be assigned to an egress forwarding plane of an HSQ IOM. The policy contains the required parameters to create and size root and mid-tier buffer pools on an HSQ IOM, and apply a slope policy to each. The HS pool policy can be applied using the hs-pool-policy command within the config>card>fp <i>fp-number</i> egress context.

The system supports 63 HS pool policies including the default HS pool policy.

- **HSQ IOM System Reserved Buffers** — The HSQ IOM maintains two types of queues; provisioned queues and system reserved queues. The HSQ IOM ensures that provisioned queues cannot consume buffers that must be available for internal system queues required for correct operating behavior. To prevent buffer starvation between the two types of queues, the system divides the available buffers into two portions. The first portion is given to system root pools and is allocated to HSQ queues reserved for internal functions. The second portion is given to the provisioned or user-defined root pools and is available for egress service queues, network queues, queue-group queues or subscriber queues. By default, 5% of the total buffers available are given to the system root pools leaving 95% for the provisioned root pools.

The default separation between the system and the provisioned root pools can be overridden using the **system-reserve** command.

- **Root Pools** — Root pools are the buffer pools at the bottom of the buffer allocation hierarchy. Two sets of root pools exist: the system root pools and the provisioned root pools. The system root pools cannot be managed by the HS pool policy; only the total number of buffers given to the system root pools can be adjusted by using the **system-reserve** command.

The HS pool policy manages sixteen provisioned root pools defined under the **root-tier** context and specified as root-pool 1 through 16. Each root pool accepts a weight command that defines the relative quantity of buffers that are allocated to each root pool. Root pools are deactivated by defining a weight equal to 0. Root pools with a non-zero weight are sized based on the pool's weight divided by the sum of all root pool weights, multiplied by the available buffer space. In this manner, all buffers not reserved for system use are distributed between the provisioned root pools without oversubscription. The lack of oversubscription prevents buffer starvation between the root pools allowing root pools, to act as protected buffer space between different types of traffic (best-effort, expedited, or real-time). Root pools allocate buffers to the FP-level mid-tier pools.

- **Mid-Tier Pools** — Mid-tier pools are the buffer pools that act as aggregators for port-class pools. Multiple mid-tier pools can be mapped to a single root pool and each mid-tier pool is assigned a percentage of that root pool's buffer space. The sum of the percentages may exceed 100%, allowing for oversubscription of the root pool's buffer space. Due to statistical multiplexing principles, oversubscribing the root pool's buffer allocation may allow more efficient use of the available buffers as not all mid-tier pools are expected to use their fair share simultaneously. Examples of a multiple mid-tier pool application are multiple assured forwarding (AF) or best-effort classes being grouped together in the same root pool. The HS pool policy manages the sixteen mid-tier pools on an HSQ IOM, defined under the **mid-tier** context and specified as mid-pool 1 through 16.

- **Port Class Pools** — The HSQ IOM maintains two sets of scheduler class pools per port: a standard (or default) set and an alternate set. Each set contains six pools, one for each scheduler class serviced by the HSQ IOM port scheduler. Each queue or WRR group is mapped to a scheduling class based on the HS attachment policy defined within the policy or template used to create the queue. Within the SAP egress policy, network queue policy and egress queue group template, an **alt-port-class-pool** command specifies whether the queues created through the policy use the standard or alternate set of port class pools on the physical port. Further, the scheduling class servicing the queue defines which port-class pool (1 through 6) within the set allocate buffers to the queue.

Port-class pools are defined within the **hs-port-pool-policy**, which is applied to each physical port. Further information on HSQ IOM port-class pools is contained in the HS port pool policy section.

- **HSQ Stable Pool Sizing Equivalency** — Stable pool sizing is a feature supported on ingress and on non-HSQ egress forwarding planes. By default, the system tries to make all buffers available to active ports (provisioned and equipped). This leads to a condition where an IOM may have only a single MDA populated and the users on that MDA receive all available buffers. At a later date, the second MDA can be populated, causing the buffer space to be fragmented between the users on each MDA. The users on the earlier populated MDA may perceive a degradation in service based on the change in available buffers. The stable pool sizing feature mitigates this potential issue by segregating the buffer space per MDA.

The HSQ IOM can be made to operate in this stable buffer allocation mechanism by utilizing per-MDA buffer pools. This is accomplished by performing the following steps:

1. Create two sets of root-pools and two sets of mid-pools in the **hs-pool-policy** applied to the IOM's FP egress CLI context. The first set of mid-pools should be parented to the first set of root-pools. The second set of mid-pools should be parented to the second set of root-pools.
 2. Create two distinct HS port pool policies. One is applied to the ports on the first MDA and has the port-class pools parented to the first set of mid-tier pools from the FP level policy. The second HS port pool policy is applied to the ports on the second MDA (when it is provisioned) and has the port-class pools parented to the second set of mid-tier pools from the FP level policy. This provides deterministic pool sizing independent of MDA equipping events.
 3. Configure further control at the port-class level by utilizing **explicit-percent** based port-class pool sizing, which eliminates the effect of changing port states, including bandwidth changes.
- **HSQ Queue Buffer Allocation** — As each packet arrives at an HSQ queue, the queue must obtain buffers to admit the packet on the queue. The queue first checks the depth of the queue relative to the packets congestion priority (based on the in, out, or exceed profile) to determine if the packet should be discarded based on early congestion detection or based on the MBS threshold. If the packet is allowed into the queue, the HSQ IOM continues to determine buffer availability using checks to the queue's port-class pool, the port-class pool's mid-tier pool, and the mid-tier pool's root pool. The same RED slope type used at the queue (high, low, or exceed) is used within each buffer pool. If a buffer is available, the buffer can be allocated, and given to the queue.
 - **Default HSQ Pool Policy** — An HSQ pool policy with the name **default** always exists on the system and does not need to be created. The default pool policy cannot be changed and is used by all HSQ IOMs within the system unless an explicitly created **hs-pool-policy** is associated with a forwarding plane.

The default policy contains the following parameters:

system-reserve: 5%

root-pool 1

allocation-weight:75

slope-policy:_tmnx_hs_default

root-pool 2

```
allocation-weight:25
slope-policy:_tmnx_hs_default
root-pool 3 to 16
allocation-weight:0
slope-policy:_tmnx_hs_default
mid-pool 1
parent-root-pool:1
allocation-percent:40%
slope-policy:_tmnx_hs_default
mid-pool 2
parent-root-pool:1
allocation-percent:35%
slope-policy:_tmnx_hs_default
mid-pool 3
parent-root-pool:1
allocation-percent:30%
slope-policy:_tmnx_hs_default
mid-pool 4
parent-root-pool:1
allocation-percent:25%
slope-policy:_tmnx_hs_default
mid-pool 5
parent-root-pool:2
allocation-percent:80%
slope-policy:_tmnx_hs_default
mid-pool 6
parent-root-pool:2
allocation-percent:20%
slope-policy:_tmnx_hs_default
mid-pool 7 to 16
parent-root-pool:None
allocation-percent:1%
slope-policy:_tmnx_hs_default
```

The **no** form of the command removes the HS pool policy from the system. If the HS pool policy is currently associated with a forwarding plane, the command fails.

Parameters *policy-name* — Specifies pool policy name, up to 32 characters. Each HS pool policy must be uniquely named within the system.

create — This keyword is required when first creating the configuration context. After the context is created, it is possible to navigate into the context without the **create** keyword.

mid-tier

Syntax	mid-tier
Context	config>qos>hs-pool-policy
Description	This command enables the context to configure HS pool policy parameters. Within the mid-tier context, mid-pools can be associated with a root pool, sized as a percentage of the root pool, have an HS slope policy applied, or be configured with a port bandwidth oversubscription factor parameter used to influence the port-class pool sizes associated with the mid-tier pool.

mid-pool

Syntax	[no] mid-pool <i>mid-pool-id</i>
Context	config>qos>hs-pool-policy>mid-tier
Description	This command enables the context to configure mid-pool tier parameters for an HS pool policy. Parameters allow for allocating the percentage of the root pool size, defining a mid-tier pool's root-pool parent, specifying the port bandwidth oversubscription factor, or specifying a slope policy for the specific mid-tier pool. The no form of the command reverts the parent root pool association to root-pool 1, reverts to the default allocation-percentage value, the default port-bw-oversub-factor, and default slope-policy to the specified mid-pool.
Parameters	<i>mid-pool-id</i> — Specifies the mid pool ID. This is a required parameter when this command is executed and specifies which mid-pool context is being entered.
Values	1 to 16

allocation-percent

Syntax	allocation-percent <i>percent-of-parent-pool</i> no allocation-percent
Context	config>qos>hs-pool-policy>mid-tier>mid-pool
Description	This command sizes the associated mid-pool based on the specified percent of the parent pool. The size is obtained by applying the specified percentage value to the current root-pool size acting as the mid-pool's parent. Whenever the parent root-pool is changed to a new root-pool or the size of the current parent root-pool is modified, the mid-pool's size is updated.

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

Default	allocation-percent 1.00
Parameters	<i>percent-of-parent-pool</i> — Specifies the percent of the parent pool. This parameter is required when the allocation-percent command is executed. This parameter defines the percentage of the root pool's size to derive the size of the mid-pool. The value is specified as a percentage with two decimal places (100th of a percent).
Values	0.01 to 100.00

parent-root-pool

Syntax	parent-root-pool <i>root-pool-id</i> no parent-root-pool
Context	config>qos>hs-pool-policy>mid-tier>mid-pool
Description	This command creates a buffer allocation mapping between the associated mid-pool <i>mid-pool-id</i> and the specified parent-root-pool <i>root-pool-id</i> . The specified root pool ID must have a non-zero allocation-weight or the command fails. After a mid-pool is successfully associated with a root-pool, the parent root-pool's allocation-weight value cannot be set to zero. When the <i>root-pool-id</i> is set to none , no buffers are assigned to the mid-tier pool. The no form of the command reverts to the default.
Default	parent-root-pool 1
Parameters	<i>root-pool-id</i> — Defines the parent root pool to which the mid-pool is associated. This parameter is required when executing the parent-root-pool command.
Values	1 to 16, none

port-bw-oversub-factor

Syntax	port-bw-oversub-factor <i>oversubscription-factor</i> no port-bw-oversub-factor
Context	config>qos>hs-pool-policy>mid-tier>mid-pool
Description	This command modifies the size of the mid-pool when calculating the port-class pool sizes based on port bandwidth ratios. The command does not actually change the size of the mid-pool, only the size reported to the port-class pool sizing function.

Port-class pools can be sized in one of two ways: dynamically (proportionate to the bandwidth of each port) or explicitly (based on a percentage of the parent mid-pool). Explicit percentages require careful determination of the amount to give each pool. The dynamic sizing function attempts to automatically size each pool based on the relative amount of bandwidth each port-class pool is supporting compared to other port's port-class pools. This is accomplished by determining a dynamic weight for each port with port-class pools mapped to a given mid-pool. As true with any weighted behavior, the mid-pool buffer allocation resource is distributed in a non-oversubscribed manner to its child port-class pools. The **port-bw-oversub-factor** *oversubscription-factor* allows this distribution mechanism to become proportionally oversubscribed based on the defined factor. An oversubscription-factor of 1.5 causes the port-class pool dynamic sizes to be 1.5 times bigger, allowing for a potentially more efficient utilization of the buffers represented by mid-pool.

The **port-bw-oversub-factor** *oversubscription-factor* for a mid-pool can be modified at any time, causing the corresponding port-class pool dynamic sizes to be recalculated.

A similar behavior can be obtained by increasing the mid-pool's allocation-percent of its parent root-pool. However, the major difference in using **port-bw-oversub-factor** is that it provides larger port-class pools without allowing the mid-pool to use a higher number of buffers in the root pool.

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

Default	port-bw-oversub-factor 1
Parameters	<i>oversubscription-factor</i> — Specifies the factor by which the dynamically-sized port-class pools associated with the mid-pool may oversubscribe the mid-pool. This parameter is required when the port-bw-oversub-factor command is executed.
Values	1 to 10

slope-policy

Syntax	slope-policy <i>policy-name</i> no slope-policy
Context	config>qos>hs-pool-policy>mid-tier>mid-pool config>qos>hs-pool-policy>root-tier>root-pool
Description	This command specifies the slope policy to be used to define the high, low, and exceed slopes within the pool. The slope (high, low, or exceed) used on the egress queue for the packet that generated the buffer request is also used in the mid-pool from the applied slope policy. The pool's current allocation amount is applied to the appropriate slope to derive the buffer rejection probability. The probability value is compared to a randomly-generated number. If the probability decision generates a rejection decision or the buffer pool has no remaining free buffers, the buffer request fails and the arriving packet is discarded. Otherwise, a buffer is allocated as long as the port-class and root-tier buffer pools also honor the buffer request.

The **no** form of the command restores the default slope policy to the associated pool.

Default	slope-policy _tmnx_hs_default
Parameters	<i>policy-name</i> — Specifies the slope policy associated with this pool, up to 32 characters.

root-tier

Syntax	root-tier
Context	config>qos>hs-pool-policy
Description	This command enables the context to configure root pool parameters. Within the root-tier context, root pools can be sized using the allocation-weight command or a slope policy can be associated with a root pool.

root-pool

Syntax	[no] root-pool <i>root-pool-id</i>
Context	config>qos>hs-pool-policy>root-tier
Description	This command enables the context to configure root tier parameters. Within the root-tier context, root pools can be sized using the allocation-weight command or a slope policy can be associated with a root pool. The no form of the command restores the default allocation-weight value and default slope policy to the specified root pool. Root pool 1 has a different default weight than root pools 2 through 16. The no root-pool command fails for root pools 2 through 16 if the root pool is currently the parent of a mid-tier pool.
Parameters	<i>root-pool-id</i> — Specifies the root pool ID. This is a required parameter when executing the root-pool command and specifies which root pool context is being entered.
Values	1 to 16

allocation-weight

Syntax	allocation-weight <i>pool-weight</i> no allocation-weight
Context	config>qos>hs-pool-policy>root-tier>root-pool
Description	This command specifies the weight that is applied to the root pool and is divided by the sum of all root pool weights to derive the pool's buffer allocation factor. The amount of buffers remaining after the system-reserve percentage is applied is multiplied by the buffer allocation factor to derive the pool size.

Root pools function as an oversubscription control mechanism. A root pool acts as the root of a hierarchy of buffer pools and queues with respect to buffer allocation. Because the sum of the root pool sizes does not exceed the total number of buffers available, the number of buffers indicated by the root pools size is always be available to the queues within the root pools hierarchy, queues from one hierarchy can never steal buffers from another.

A root pool hierarchy is based on the dynamic parenting of one or more mid-tier pools to a root pool. A mid-tier pool represents the buffering allowed for all port-class pools mapped to the mid-tier pool. Each mid-tier pool is sized as a percentage of the root pool to which it is parented. The sum of the mid-tier pools percentages for a root pool may be greater than 100 percent, which allows the root pool to be oversubscribed. This can be beneficial when large fluctuations in mid-tier buffer utilization are expected and a given mid-tier pool should be allowed to exceed its fair share of buffering.

Through the mapping hierarchy presented above, each queue is mapped to a port-class pool, mid-tier pool, and root pool.

A root pool with an **allocation-weight** set to “0” is considered inactive and is not allocated buffers. Mid-tier pools cannot be parented to a root pool with a weight set to “0”. After a mid-tier pool is associated with a root pool, the root pool’s weight cannot be set to “0”.

As port classes are mapped to mid-tier pools in a different policy than mid-tier pools are mapped to root pools, a port-class pool can be mapped to a mid-tier pool that is not parented to a root pool. A queue mapped indirectly to a non-parented mid-tier pool has its operational MBS value set to zero and drops all incoming packets.

When a root pool’s allocation weight is modified, all root pools, mid-tier pools, and port class pool sizes are reevaluated and modified when necessary.

The **no** form of the command restores the default **allocation-weight** value to the associated root pool. Root pool 1 has a different default weight than root pools 2 through 8. The **no allocation-weight** command fails for root pools 2 through 8 if the root pool is currently parented to a class pool.

Default	root-pool 1: 100 root-pool 2 to 16: 0
Parameters	<i>pool-weight</i> — Defines the weight of the associated root-pool <i>root-pool-id</i> and is used by the system to calculate the size of the root buffer pool. This parameter is required when executing the allocation-weight command. Setting the <i>pool-weight</i> to 0 disables the pool and prevents the root pool from being a parent to any class pools. Root pool 1 cannot be set with an allocation weight of 0.
Values	root-pool 1:1 to 100 root-pool 2 to 16:0 to 100

system-reserve

Syntax	system-reserve <i>percent-of-buffers</i> no system-reserve
Context	config>qos>hs-pool-policy
Description	This command defines the amount of HSQ IOM buffers that is set aside for internal system use. By default, 5% of the total buffer space is reserved for system internal queues. The command is provided for the case where the reserved buffer space is either insufficient or excessive. Exercise care when modifying this value. When the system reserve value is changed, all the provisioned port-class, mid-tier, and root pool sizes are reevaluated and possibly changed. Use the show hs-pools card-slot-number fp forwarding-plane egress command to display the current buffer allocation and buffer usage conditions on an HSQ IOM. The no form of the command reverts to the default system reserve value.
Default	system-reserve 5.0
Parameters	<i>percent-of-buffers</i> — Specifies the percentage of HS buffers that are reserved for internal system use. This parameter is required when executing the system-reserve command. The parameter accepts a percent value with two decimal places (100th of a percent). Values 1.00 to 30.00

15.4.2.4 HS Port Pool Policy Commands

hs-port-pool-policy

Syntax	hs-port-pool-policy <i>policy-name</i> [create] no hs-port-pool-policy <i>policy-name</i>
Context	config>qos
Description	This command creates an HS port buffer pool policy. The policy can be assigned to an egress port on an HSQ IOM. The policy contains the required commands to define and size port-class buffer pools on an HSQ IOM. The policy can be applied using the hs-port-pool-policy command within the config>port>ethernet>egress context. SR OS supports 2047 HS port pool policies including the default HS port pool policy.

HSQ IOM port buffer pools provide buffer control for queues based on the queue's scheduling class. Two sets of scheduling class pools exist per port: a standard (or default) set and an alternative set. The SAP egress policy, network queue policy, and egress queue group template have a parameter (**alt-port-class-pool**) that specifies that the queues created by the policy or template uses the alternate port-class pools, as opposed to the default standard port-class pools. Each set has six pools, one for each scheduling class. Based on the **alt-port-class-pool** setting and the queue's scheduling class (based on the HS attachment policy configuration), each queue is mapped to a specific port-class pool.

The HS port pool policy defines how each of these pools are parented (mapped) to an FP level mid-pool and how each pool is sized.

The system allows two separate mechanisms to size each port-class pool:

- dynamic sizing based on port bandwidth, relative to bandwidth of other ports
- explicit sizing based on a percentage of the port-class pool's parent mid-pool

Dynamic Port-Class Pool Sizing — Dynamic port-class pool sizing is a mechanism that provides a fair share of a mid-pool's size to each of the port-class pools based on the potential bandwidth represented by each port. To understand port-class pool sizing, consider the following:

- Each port's bandwidth is the minimum of the port's line rate, the port's configured **egress-rate**, and the port's **hs-scheduler-policy max-rate**.
- The port's bandwidth can be further modified by the port's **egr-percentage-of-rate** command, which increases or decreases the port's bandwidth derived by the specified percent. This parameter allows the port to have a higher or lower bandwidth-derived weight based on how the port is actually being used instead of bandwidth alone.
- Because the port-class pools are user mapped to the mid-pools, not every port has a port-class pool associated with a mid-pool, requiring that the system perform the relative bandwidth calculations separately per mid-pool.
- Each port's portion of a specified mid-pool's size is calculated based on:

$$\text{Port\_Portion} = (\text{Port\_Adj\_Bw} / \text{Sigma\_Mid\_Pool\_Ports\_Adj\_Bw}) * \text{Mid\_Pool\_Size}$$
 where $\text{Sigma\_Mid\_Pool\_Ports\_Adj\_Bw}$ is the sum of the adjusted bandwidths for all ports with port-class pools mapped to the mid-pool that are not sized using explicit-percent.
- A port without any port-class pools associated with a given mid-pool has a port portion of zero for that mid-pool.
- Multiple port-class pools on the same port can be mapped to the same mid-pool, requiring a mechanism to distribute the portion of the mid-pool given to the port between multiple port-class pools. Each mid-pool's **port-bw-weight** parameter is used to determine how much of the port's mid-pool portion is given to each port-class pool associated with mid-pool. Port-class pools sized using an **explicit-percent** value instead of port bandwidth are assumed to have a **port-bw-weight** equal to 0, causing those port-class pools to not participate in the port portion distribution. It is expected (but not required) that one of port bandwidth-based sizing or explicit percent-based sizing is used and any concurrent use of both mechanisms is transitory in nature.

- The port bandwidth weighting mechanism allocates 100% of the mid-pool size to the associated port-class pools. To allow the port-class pools to oversubscribe the parent mid-pool, a mid-pool **port-bw-oversub-factor** parameter is supported that allows the port-class pools sized by dynamic port bandwidth to increase in size by the specified oversubscription factor. This oversubscription factor can provide a more efficient use of the mid-pool's available buffers because it is not expected that all port-class pools are utilizing their allotted size simultaneously.

Explicit Port-Class Pool Sizing — The port-class pool's allocation **explicit-percent** *percent-of-parent-pool* command is used to override the dynamic pool sizing mechanism for a given mid-pool. The specified percentage value is applied to the port-class's parent mid-pool's size to derive the port-class pool size.

Explicit and Dynamic Sizing from the Same Mid-Pool Parent — If explicit and dynamic pool sizing are used simultaneously for port-class pools parented to the same mid-pool, unexpected contention or under utilization of the mid-pool's available buffers may result. While this is not a proscribed condition, it is expected most instances of dual-sizing mechanisms are transitory, based on moving between the two mechanisms.

Port-Class Pool Slope Policy Association — The HS port pool policy also provides the ability to specify a slope policy on each port-class pool. The slope policy is used to define the high, low, and exceed slope parameters used to manage contention within the port-class pool.

Default HS Port Pool Policy — An HS port pool policy with the name **default** always exists on the system and does not need to be created. The default port pool policy cannot be changed and is used by all HSQ IOMs within the system unless an explicitly created **hs-port-pool-policy** is associated with an HSQ egress port.

The default policy contains the following parameters:

Standard Port-Class Pools

Port-Class-Pool 1

Parent:Mid-Pool 1

Port-Class-Pool 2

Parent:Mid-Pool 2

Port-Class-Pool 3

Parent:Mid-Pool 3

Port-Class-Pool 4

Parent:Mid-Pool 4

Port-Class-Pool 5

Parent:Mid-Pool 5

Port-Class-Pool 6
 Parent:Mid-Pool 6
 Port-Class-Pool 1 to 6
 Port-Bw-Weight:1
 Slope-Policy:_tmnx_hs_default
 Alternate Port-Class Pools
 Port-Class-Pool 1 to 6
 Parent:None
 Port-Bw-Weight:1
 Slope-Policy:_tmnx_hs_default

Parameters *policy-name* — Specifies an HS port pool policy name up to 32 characters.
create — This keyword is required when first creating the configuration context. After the context is created, it is possible to navigate into the context without the **create** keyword.

alt-port-class-pools

Syntax **alt-port-class-pools**
Context config>qos>hs-port-pool-policy
Description This command enables the context to configure alternate port class pools parameters. Within this context, the corresponding port-class pools can be associated with a mid-pool, explicitly sized as a percentage of the mid-pool size, dynamically sized based on relative port bandwidth, or have a slope policy applied.

std-port-class-pools

Syntax **std-port-class-pools**
Context config>qos>hs-port-pool-policy
Description This command enables the context to configure standard port-class pools parameters. Within this context, the corresponding port-class pools can be associated with a mid-pool, explicitly sized as a percentage of the mid-pool size, dynamically-sized based on relative port bandwidth, or have a slope policy applied.

class-pool

Syntax	[no] class-pool <i>alt-class-pool-id</i>
Context	config>qos>hs-port-pool-policy>alt-port-class-pools
Description	This command enables the context to configure a class pool's parent mid-pool, dynamic port bandwidth weight, explicit percentage of mid-pool size, or a slope policy. Six alternate port-class pools always exist (one for each of the six scheduling classes) and do not need to be created. The no form of the command restores the default parent-mid-pool association to mid-pool none, restores the default allocation port-bw-weight 1 setting (explicit-percent disabled), and restores the default slope policy to the specified class-pool.
Parameters	<i>alt-class-pool-id</i> — Specifies the class pool ID.
Values	1 to 6

class-pool

Syntax	[no] class-pool <i>std-class-pool-id</i>
Context	config>qos>hs-port-pool-policy>std-port-class-pools
Description	This command enables the context to configure class pool's parent mid-pool, dynamic port bandwidth weight, explicit percentage of mid-pool size, or a slope policy. Six alternate port-class pools always exist (one for each of the six scheduling classes) and do not need to be created. The no form of the command restores the default parent-mid-pool association to mid-pool 1, restores the default allocation port-bw-weight 1 setting (explicit-percent disabled), and restore the default slope policy to the specified class-pool.
Parameters	<i>std-class-pool-id</i> — Specifies the class pool ID.
Values	1 to 6

allocation

Syntax	allocation explicit-percent <i>percent-of-parent-pool</i> allocation port-bw-weight <i>pool-weight</i> no allocation
Context	config>qos>hs-port-pool-policy>alt-port-class-pools>class-pool config>qos>hs-port-pool-policy>std-port-class-pools>class-pool

Description	This command sizes the associated class-pool based on either the specified explicit-percent <i>percent-of-parent-pool</i> or based on the dynamic port bandwidth portioning mechanism. Setting an explicit percentage prevents the port-class pool from participating in the dynamic port level bandwidth based distribution of the mid-pool's size as the port bandwidth weight of the port-class pool becomes zero (0). Setting a port bandwidth weight causes the explicit percent value to become zero (0) disabling explicit sizing of the port-class pool. The no form of the command sets the <i>percent-of-parent-pool</i> value to zero (0) and the <i>pool-weight</i> parameter to 1 for the port-class pool, restoring the default settings.
Default	allocation 1
Parameters	<i>percent-of-parent-pool</i> — Specifies the percentage of parent pool being allocated. This parameter must be configured when specifying the explicit-percent. The <i>percent-of-parent-pool</i> value is expressed as a percentage with two decimal places (100th of a percent) that indicates that the port-class pool should be sized by applying the value to the parent mid-pool size. Specifying explicit-percent forces the port-bw-weight to a zero (0) value (disabled). Values 0.01 to 100.00 <i>pool-weight</i> — Specifies port bandwidth weight being allocated. The port-bw-weight and explicit-percent commands are mutually exclusive. The <i>pool-weight</i> parameter is required when specifying the port bandwidth weight and defines both that the port-class pool should be sized in the port bandwidth distribution of the mid-pool's size and what the distribution weight should be for the port-class pool compared to other port-class pools associated with the same mid-pool when competing for the port's distribution portion. Values 1 to 100

parent-mid-pool

Syntax	parent-mid-pool <i>mid-pool-id</i> no parent-mid-pool
Context	config>qos>hs-port-pool-policy>alt-port-class-pools>class-pool config>qos>hs-port-pool-policy>std-port-class-pools>class-pool
Description	This command creates the buffer allocation mapping between the associated class pool and the specified mid-pool. Use care when selecting a mid-pool in an active state (properly mapped to a root-pool with a non-zero allocation percentage). If a port-class pool is parented by an inactive mid-pool, the queues using the port-class pool are forced into an operational MBS setting of 0, causing all packet to be discarded. A port-class pool can be made inactive (no available buffers) by executing parent-mid-pool none in the port-class pool context. The no form of the command reverts to the class-pool parenting value. For the standard port-class pools, this default is 1. For alternate port-class pools the default is none.

Default	alt-port-class-pools: none std-port-class-pools: 1
Parameters	<i>mid-pool-id</i> — Specifies the mid-pool identifier in the HS pool policy. Either a valid mid-pool ID or none must be specified when executing the parent-mid-pool command. The <i>mid-pool-id</i> parameter defines the parent mid-pool to which the port-class is associated. The none keyword deactivates the port-class pool, causing the pool to have a zero size. A queue can still map to an inactive port-class pool although all packets are discarded by the queue. Values 1 to 16, none

slope-policy

Syntax	slope-policy <i>policy-name</i> no slope-policy
Context	config>qos>hs-port-pool-policy>alt-port-class-pools>class-pool config>qos>hs-port-pool-policy>std-port-class-pools>class-pool
Description	This command specifies the slope policy that is used to define the high, low, and exceed slopes within the port-class pool. The slope used on the egress queue for the packet that generated the buffer request is also used in the class-pool. The pool's current allocation amount is applied to the appropriate slope to derive the buffer rejection probability. The probability value is compared to a randomly generated number. If the probability decision generates a rejection or the buffer pool has no remaining free buffers, the buffer request fails and the arriving packet is discarded. Otherwise, a buffer is allocated as long as the mid-tier and root-tier buffer pools also honor the buffer request. The no form of the command restores the default slope policy to the associated class-pool.
Default	slope-policy _tmnx_hs_default
Parameters	<i>policy-name</i> — Specifies the policy name, up to 32 characters. This parameter is required when executing the slope-policy command and must refer to an existing slope policy within the system. If a slope policy with the specified name does not exist, the slope-policy command fails without modifying the slope behavior on the pool. A non-existent slope-policy error is generated. After a slope policy is associated with any HSQ queue or buffer pool, the policy cannot be deleted.

15.4.2.5 HS Scheduler Policy Commands

hs-scheduler-policy

Syntax	hs-scheduler-policy <i>policy-name</i> [create] no hs-scheduler-policy <i>policy-name</i>
Context	config>qos
Description	This command configures an HS scheduler policy. The HS scheduler policies are applied to egress HSQ ports in the config>port>ethernet>egress context. The policy contains the required commands to provision the scheduling behavior of the HSQ scheduler classes. When assigned to an HSQ egress port, the policy is used to define the scheduling behavior for all queues associated with the egress port. The values defined in the policy can be overridden on each scheduler instance using the config>port>ethernet>egress>hs-scheduler-overrides command.

HSQ Queue Groups — A fundamental concept on an HSQ IOM is the queue group. Queue groups are not directly managed by the provisioning. Instead, they are indirectly assigned when creating SAPs or subscribers on an HSQ port. A queue group has eight queue members, numbered from 1 through 8. When creating a SAP or subscriber associated on an HSQ egress port, a queue group is allocated to the object. Within the SAP egress policy, network queue policy and egress queue group template, provisioned queue IDs 1 through 8 correspond directly to queue group queue IDs 1 through 8. Each group also allows each queue to be dynamically placed in a scheduling class or on one of two WRR groups local to the queue group.

Each queue within the group has three RED slopes (managed by associating a slope policy to the queue), an MBS defined in bytes, a packet byte offset parameter used to add or subtract bytes to or from each packet handled by the queue for accounting purposes, and a PIR shaper used to rate limit the queue. An HSQ attachment policy associated with the queue group defines how each queue maps either directly to a scheduling class or one of the WRR groups within the queue group. The attachment policy also defines the scheduling class attachments for the WRR groups.

The queue group supports an aggregate shaper used to manage an aggregate rate limit for all queues within the group. Scheduling for queues within the queue group is stopped and started based on the rate set on the shaper.

Scheduling Classes and Scheduling Priorities — HSQ supports six scheduler classes (1 through 6). The scheduler class should not be confused with a QoS policy forwarding class. Forwarding classes within the system are used between the ingress and egress forwarding complexes and help the system to map a packet to per-hop and per-domain behavior. Scheduling classes are slices of scheduling opportunity within a port-scheduling context. Each port scheduler maintains six strict priority levels, where 6 is the highest priority and 1 is the lowest. As a rule, the scheduler services all active queues associated with priority level 6 before moving to queues on priority level 5. This strict behavior continues through priority

level 1. Scheduling classes are mapped either to their corresponding scheduling priority level (scheduling class 1 mapped to priority level 1 through scheduling class 6 mapped to priority level 6) or to a single port level WRR group. The WRR group allows collapsing up to 6 of the scheduling classes into a single scheduling priority. The WRR group provides a weighted fair scheduling behavior for its member scheduling classes at that strict priority level.

Strict Priority Level PIR — The scheduler supports a strict scheduling level PIR that limits the amount of bandwidth allowed for the level. The rate is defined in increments of megabits per second and can be set to **max** (the default setting) which disables the shaping function. The scheduler includes the full Ethernet frame encapsulation overhead when updating the priority level PIR, including the 12-byte inter-frame gap and the 8-byte preamble.

Scheduler Maximum Rate — A maximum scheduling rate can be defined for the scheduler. The rate is specified in megabits per second and the default rate is **max** which allows the scheduler to operate without a set limit. When the HS scheduling policy is applied to an egress port, the maximum scheduling rate can be used to define a rate less than the available line rate of the port. The scheduler includes the full Ethernet frame encapsulation overhead when updating the scheduler level PIR, including the 12 byte inter-frame gap and the 8-byte preamble.

HS Scheduler Policy Overrides — After an HS scheduler is applied to an egress port, the various parameters can be overridden, allowing an HS scheduler policy to be adapted to changing needs on a port without requiring a new policy to be created.

Default HS Scheduling Policy — An HS scheduling policy with the name **default** always exists on the system and does not need to be created. The default policy cannot be modified or deleted.

The default policy contains the following parameters:

Table 120 HS Scheduler Policy Parameter Defaults

Parameter	Sub-Parameter	Default
max-rate		max
scheduling-class 1 through 6	rate	max
	group	N/A
	weight	N/A
group 1	rate	max

Default The **no** form of the command removes an HS scheduler policy from the system. If the HS scheduler policy is currently associated with an egress port, the command fails.

Parameters *policy-name* — Specifies the HS scheduler policy up to 32 characters. Each HS scheduler policy must be uniquely named within the system.

create — This keyword is required when first creating the configuration context. After the context is created, it is possible to navigate into the context without the **create** keyword.

group

Syntax	group <i>group-id</i> rate <i>rate</i> no group <i>group-id</i>
Context	config>qos>hs-scheduler-policy
Description	This command defines the maximum rate allowed for the scheduling classes mapped to the specified <i>group-id</i>. A group is a scheduling component used to combine up to six consecutive scheduling classes into a single strict priority level. Each scheduling class within the group has an associated weight. When the scheduler is servicing the strict level associated with the group, the ratio of bandwidth allocated to each scheduling class within the group during congestion is relative to the ratio of the weight of each active member. The no form of the command reverts to the default.
Default	group 1 rate max
Parameters	<i>group-id</i> — Specifies the group ID. The group always exists and does not need to be created prior to defining group membership. Values 1 <i>rate</i> — Specifies the maximum rate in megabits per second. When the max keyword follows the rate keyword, the bandwidth limitation is removed from the group. The max keyword and the <i>rate</i> parameter are mutually exclusive. Either max or a rate value must follow the rate keyword. Values 1 to 100000, max

max-rate

Syntax	max-rate <i>rate</i> no max-rate
Context	config>qos>hs-scheduler-policy
Description	This command defines an explicit maximum frame based bandwidth limit for the HS scheduler policy scheduler context. If a maximum rate is defined that is smaller than the port rate, the port is rate-limited to the configured megabits per second value. This command can be executed at any time for any non-default existing HS scheduler policy.

The **no** form of the command removes an explicit rate value from the HS scheduler policy. After the explicit rate value is removed, all instances of the scheduler policy on HSQ egress ports are allowed to run at the available line rate unless the instance has a max rate override in place.

Parameters *rate* — Specifies the explicit maximum frame based bandwidth limit, in megabits per second. This parameter is required when executing this command.

Values 1 to 100000, max

scheduling-class

Syntax **scheduling-class** *class-id* **group** *group-id* [**weight** *weight-in-group*]
scheduling-class *class-id* **rate** *rate*
no scheduling-class *class-id*

Context config>qos>hs-scheduler-policy

Description This command configures the behavior of a specific scheduling class on all HSQ schedulers associated with the policy. The **scheduler-class** command performs one of two operations: it configures a maximum rate for the scheduling class or places the scheduling class into the weighted scheduling group. The two operations are mutually exclusive.

By default, none of the scheduling classes are members of the weighted scheduling group and each class is set to a rate limit of **max** (no rate limit applied).

Specifying Scheduling Class Rate (or Removing the Scheduling Class from Group) — If the **scheduling-class** command is executed with the **rate** keyword specified, either **max** or a specified rate value must follow. If a *class-id* was previously mapped into the weighted scheduling group, the class is removed from the group. However, if removing the class from the group causes the group to no longer have contiguous class members, the command fails with no effect on the specified class. A “non-contiguous grouping error” is returned. The lowest or highest members within a weighted group must be removed prior to removing the middle members. For example, if scheduling classes 3, 4, and 5 were members of weighted group 1, class 4 cannot be removed first.

This command using the **rate** keyword also fails when an override for the group weight is in place on the scheduling class within a scheduler associated with the policy. The override expects the class to be associated with a weighted scheduling group and the policy rate definition is attempting to remove the class from the group. An “override mismatch” error is generated, specifying the scheduling object where the override exists.

After a rate has been successfully defined for a scheduling class, the specified rate is automatically updated on all HSQ scheduler instances associated with the scheduling policy. The exception is where the scheduler instance has a local override for the rate on the scheduling class.

Specifying Scheduling Class Weighted Group Membership — If the **scheduling-class** command is executed with the **group** keyword specified, the group ID value of 1 must follow. The corresponding optional **weight** keyword is used to specify the weight of the scheduling class within the group. If weight is not specified, the default weight of 1 is used. If the specified scheduling class is not contiguous with the other scheduling classes in the group, the command fails with no change to the current state of the scheduling class and a “non-contiguous grouping” error is returned, specifying the weighted scheduling group and the current group members.

The **scheduling-class** command fails using the **group** keyword when a rate override for the scheduling class exists on an HSQ scheduler instance associated with the policy. The rate override for the scheduling class indicates the class is directly attached to a strict priority level, conflicting with the policy **group** keyword trying to place the class in the specified group. The command fails without affecting the scheduling class definition on the policy and returns an error specifying the scheduling object where the override exists.

Other Override Constraints — The scheduling overrides cannot change or remove a scheduling class from a policy-defined weighted group membership.

The **no** form of the command returns the scheduling class represented by *class-id* to the default behavior. The default behavior for a scheduling class is to not be a member of the weighted scheduling class group and have a rate set to **max**. The **no scheduling-class** command fails if the scheduling class is currently a member of the weighted scheduling class group and a weight override is in effect on a scheduling object for the class, in which case an error is returned.

Parameters *class-id* — Specifies the scheduling class for HS scheduler policy. The *class-id* value is a required parameter that specifies which scheduling class the scheduling-class command is acting upon.

Values 1 to 6

group-id — Specifies the group this HS scheduler policy scheduling class belongs to. The **group** and the **rate** keywords are mutually exclusive when executing this command. A group ID value of 1 must follow the **group** keyword. The **group** keyword removes the class ID from its inherent strict scheduling level and places it into the specified group ID. The associated **weight** parameter is optional and is used to specify the weight of a class ID within the weighted scheduling class group. Specifying the **group** parameter while an override for the scheduling class exists for rate causes the **scheduling-class** command to fail.

Values 1

weight-in-group — Specifies the weight the HS scheduler policy should apply to this scheduling class within the group in which it belongs. This keyword is optional and must follow the **group** parameter when specified. The *weight-in-group* parameter must follow the **weight** keyword and is used to specify the relative weight of the *class-id* to the other scheduling classes within the group. If the group is specified without the **weight** parameter, a default weight of 1 is used.

Values 1 to 127

rate — Specifies the explicit maximum frame-based bandwidth limit, in megabits per second, for this HS scheduler policy scheduling class. The **rate** and **group** keywords are mutually exclusive. Either the **rate** or **group** keyword must be specified when executing this command. When specified, the **rate** keyword must be followed by either the keyword **max** or a rate specified in megabits per second. The specified rate can be overridden at the port Ethernet egress using the scheduler override functions. A newly-created HS scheduler policy defaults each scheduling class to have its rate set to **max** and the weighted scheduling class group has no members.

Values 1 to 100000, max

The **max** keyword specifies that a limit is not enforced for the specified class ID and that the class ID is not a member of a weighted scheduling class group. The **max** keyword and the *rate* value are mutually exclusive; when **max** is specified, it must directly follow the **rate** keyword. Setting the rate of the class fails when the class currently has a group weight override defined on a scheduling object.

15.4.2.6 Network Queue HS Attachment Policy Commands

hs-attachment-policy

Syntax	hs-attachment-policy <i>policy-name</i> no hs-attachment-policy
Context	config>qos>network-queue
Description	This command associates an existing HS attachment policy with the network queue QoS policy. The HS attachment policy controls how the network queues are attached to scheduler classes or WRR groups, and how WRR groups are attached to the scheduler classes. It also defines the mapping of the scheduling classes to the queues' aggregate shaper's low and high burst limit thresholds. Only one HS attachment policy can be associated with a network queue policy. The no form of the command removes the policy name from the configuration and reapplies the default HS attachment policy.
Parameters	<i>policy-name</i> — Specifies an existing attachment policy up to 32 characters. Each HSQ attachment policy must be uniquely named within the system.

hs-wrr-group

Syntax [**no**] **hs-wrr-group** *group-id*

Context	config>qos>network-queue
Description	This command enables the context to configure HS WRR group information in the network queue policy. This command provisions the rate and class weight of each of the two WRR scheduling groups that can be utilized by the egress queue-group instance HSQ queues. The no form of the command reverts the HS WRR group parameters to their default values.
Parameters	<i>group-id</i> — Specifies the HS WRR group identifier. WRR group ID 1 or 2 must be specified when executing the hs-wrr-group command. The specified group ID identifies which WRR group context is entered for editing. Values 1, 2

adaptation-rule

Syntax	adaptation-rule [pir adaptation-rule] no adaptation-rule
Context	config>qos>network-queue>hs-wrr-group
Description	This command specifies how the system should resolve differences between the specified scheduling limit derived from the WRR group's rate command and the actual operational rate obtainable in hardware. The min, max, and closest mutually exclusive keywords specify whether the next highest rate, next lowest rate, or closest rate should be selected by the system. The no form of the command reverts to the default value.
Default	adaptation-rule pir closest
Parameters	<i>adaptation-rule</i> — Specifies the adaptation rule (min, max, or closest) to be used while computing the operational PIR value. The adaptation rule specifies the rules to compute the operational values while maintaining minimum offset. The min, max, and closest keywords are mutually exclusive. Values min — When min is specified, the queue's rate parameter is treated as the minimum rate to shape the queue. The hardware chooses the appropriate timers and PIR leaky bucket behavior to ensure that the queues shaping rate is the closest possible value without going under the specified rate. max — When max is specified, the queue's rate parameter is treated as the maximum rate to shape the queue. The hardware chooses the appropriate timers and PIR leaky bucket behavior to ensure that the queue's shaping rate is the closest possible value without going over the specified rate.

closest — When **closest** is specified, the queue's rate parameter is treated as the target rate to shape the queue. The hardware chooses the appropriate timers and PIR leaky bucket behavior to ensure that the queues shaping rate is the closest possible value and can be higher or lower than the specified rate.

hs-class-weight

Syntax	hs-class-weight <i>weight</i> no hs-class-weight
Context	config>qos>network-queue>hs-wrr-group
Description	This command specifies the class weight of this WRR group at its parent primary shaper, relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class. This allows the capacity available at the primary shaper scheduling class to be shared in a WRR manner between the HSQ queue group queues and WRR groups attached to that scheduling class. The hs-class-weight <i>weight</i> can be used to give unequal shares of the available capacity to different types of service offerings. The no form of the command reverts to weight to the default value.
Default	hs-class-weight 1
Parameters	<i>weight</i> — Specifies the class weight of the HS WRR group. Values 1, 2, 4, 8

rate

Syntax	rate <i>percent</i> no rate
Context	config>qos>network-queue>hs-wrr-group
Description	This command specifies the scheduling rate applied to the HS WRR group as a percentage of the port rate, which includes both the egress-rate and HS scheduler policy max-rate , if configured. The no form of the command reverts to the default.
Default	rate 100
Parameters	<i>percent</i> — Specifies the scheduling rate of the HS WRR group as a percentage. Values 1 to 100

hs-alt-port-class-pool

Syntax	[no] hs-alt-port-class-pool
Context	config>qos>network-queue>queue
Description	This command specifies that the HSQ queue group queues use buffers from the HS alternate port class buffer pool. The no form of the command reverts to the HSQ queue group queues using buffers from HS standard port class pools.

hs-class-weight

Syntax	hs-class-weight <i>weight</i> no hs-class-weight
Context	config>qos>network-queue>queue
Description	This command specifies the class weight of this queue at its parent primary shaper, relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class. This allows the capacity available at the primary shaper scheduling class to be shared in a WRR manner between the HSQ queue group queues and WRR groups attached to that scheduling class. The hs-class-weight <i>weight</i> parameter can be used to give unequal shares of the available capacity to different types of service offerings. This command is ignored for egress HSQ queue group queues that are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the hs-class-weight is performed under the hs-wrr-group within the network queue policy. The no form of the command reverts to the default value.
Default	hs-class-weight 1
Parameters	<i>weight</i> — Specifies class weight of the queue. Values 1, 2, 4, 8

hs-mbs

Syntax	hs-mbs <i>percent-of-queue-rate</i> no hs-mbs
Context	config>qos>network-queue>queue
Description	This command configures the queue size of an HSQ queue group network queue. Its value is calculated based on the specified percentage of one second of the queue PIR converted to bytes (the regular mbs parameter is ignored in the network queue policy).

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

Default	hs-mbs 100
Parameters	<i>percent-of-queue-rate</i> — Specifies the buffer space for the queue as a percentage of its PIR (in bytes).
Values	0.00 to 100.0

hs-wred-queue

Syntax	hs-wred-queue [<i>policy slope-policy-name</i>] no hs-wred-queue
Context	config>qos>network-queue>queue
Description	This command reverts the slope policy applied to the HSQ queue group queue to the default policy. Specifying an existing slope policy applies the named slope policy to the queue. The no form of the command reverts to the default slope policy.
Default	hs-wred-queue policy "_tmnx_hs_default"
Parameters	<i>slope-policy-name</i> — Specifies an existing slope policy to apply to this HSQ queue group queue.

hs-wrr-weight

Syntax	hs-wrr-weight <i>weight</i> no hs-wrr-weight
Context	config>qos>network-queue>queue
Description	This command specifies the WRR relative weight, with which this queue should parent into an HSQ WRR group defined within the associated HS attachment policy. The weight of each queue determines how much bandwidth that queue gets out of the total rate for the HSQ WRR group. The no form of the command reverts to the default value.
Default	hs-wrr-weight 1
Parameters	<i>weight</i> — Specifies the HS WRR group queue weight. Values 1 to 127

15.4.2.7 SAP Egress HS Attachment Policy Commands

hs-attachment-policy

Syntax	hs-attachment-policy <i>policy-name</i> no hs-attachment-policy
Context	config>qos>sap-egress
Description	This command associates an existing HS attachment policy with the SAP egress QoS policy. The HS attachment policy controls how the SAP egress queues are attached to scheduler classes or WRR groups, and how WRR groups are attached to the scheduler classes. It also defines the mapping of the scheduling classes to the queues' aggregate shaper's low and high burst limit thresholds. Only one HS attachment policy can be associated with a SAP egress policy. The no form of the command removes the policy name from the configuration and reapplies the default HS attachment policy.
Parameters	<i>policy-name</i> — Specifies an existing attachment policy up to 32 characters. Each HSQ attachment policy must be uniquely named within the system.

hs-wrr-group

Syntax	[no] hs-wrr-group <i>group-id</i>
Context	config>qos>sap-egress
Description	This command enables the context to configure HS WRR group information in the SAP egress QoS policy. The hs-wrr-group command is used to provision the rate and class weight of each of the two WRR scheduling groups that can be utilized by the SAP egress HSQ queues. The no form of the command resets the HS WRR group parameters to their default values.
Parameters	<i>group-id</i> — Specifies the HS WRR group identifier. WRR group ID 1 or 2 must be specified when executing the hs-wrr-group command. The specified group ID identifies which WRR group context is entered for editing.
Values	1, 2

adaptation-rule

Syntax	adaptation-rule [<i>pir adaptation-rule</i>] no adaptation-rule

Context	config>qos>sap-egress>hs-wrr-group
Description	This command specifies how the system resolves differences between the specified scheduling limit derived from the WRR group's rate command and the actual operational rate obtainable in hardware. The mutually exclusive min, max, and closest keywords specify whether the next highest rate, next lowest, or closest rate should be selected by the system. The no form of the command reverts to the default value.
Default	adaptation-rule pir 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. adaptation-rule — Specifies the adaptation rule (min, max, or closest) to be used while computing the operational PIR value. The adaptation rule specifies the rules to compute the operational values while maintaining minimum offset. The min, max, and closest keywords are mutually exclusive. Values min — When min is specified, the queue's rate parameter is treated as the minimum rate to shape the queue. The hardware chooses the appropriate timers and PIR leaky bucket behavior to ensure that the queues shaping rate is the closest possible value without going under the specified rate. max — When max is specified, the queue's rate parameter is treated as the maximum rate to shape the queue. The hardware chooses the appropriate timers and PIR leaky bucket behavior to ensure that the queue's shaping rate is the closest possible value without going over the specified rate. closest — When closest is specified, the queue's rate parameter is treated as the target rate to shape the queue. The hardware chooses the appropriate timers and PIR leaky bucket behavior to ensure that the queues shaping rate is the closest possible value and can be higher or lower than the specified rate.

hs-class-weight

Syntax	hs-class-weight weight no hs-class-weight
Context	config>qos>sap-egress>hs-wrr-group

Description This command specifies the class weight of this WRR group at its parent primary shaper, relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class. This allows the capacity available at the primary shaper scheduling class to be shared in a WRR manner between the HSQ queue group queues and WRR groups attached to that scheduling class. The **hs-class-weight** parameter can be used to give unequal shares of the available capacity to different types of service offerings.

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

Default hs-class-weight 1

Parameters *weight* — Specifies the class weight of the HS WRR group.

Values 1, 2, 4, 8

percent-rate

Syntax **percent-rate** *percent*
no percent-rate

Context config>qos>sap-egress>hs-wrr-group

Description This command specifies the scheduling rate applied to the HS WRR group as a percentage of the port rate, including both the port's egress rate and port's HS scheduler **policy** *max-rate*, if configured. The **percent-rate** and **rate** commands are mutually exclusive.

The **no** form of the command reverts to the rate **max**.

Parameters *percent* — Specifies the percent rate of the HS WRR group.

Values 0.01 to 100.00

rate

Syntax **rate** *rate*
no rate

Context config>qos>sap-egress>hs-wrr-group

Description This command specifies the scheduling rate applied to the HS WRR group in Kb/s. Alternatively, the keyword **max** can be specified, which removes the bandwidth limitation on the HS WRR group. The **rate** and **percent-rate** commands are mutually exclusive.

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

Default rate max

Parameters *rate* — Specifies the scheduling rate of the HS WRR group in Kb/s.

Values 1 to 2000000000, max

hs-alt-port-class-pool

Syntax	[no] hs-alt-port-class-pool
Context	config>qos>sap-egress>queue
Description	This command specifies that the HSQ queue group queues use buffers from the HS alternate port class buffer pool. The no form of the command reverts to the HSQ queue group queues using buffers from HS standard port class pools.

hs-class-weight

Syntax	hs-class-weight <i>weight</i> no hs-class-weight
Context	config>qos>sap-egress>queue
Description	This command specifies the class weight of this queue at its parent primary shaper, relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class. This allows the capacity available at the primary shaper scheduling class to be shared in a WRR manner between the HSQ queue group queues and WRR groups attached to that scheduling class. The hs-class-weight <i>weight</i> parameter can be used to give unequal shares of the available capacity to different types of service offerings. This command is ignored for egress HSQ queue group queues, which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the hs-class-weight is performed under the hs-wrr-group within the network queue policy. The no form of the command reverts to the default value.
Default	hs-class-weight 1
Parameters	<i>weight</i> — Specifies class weight of the queue. Values 1, 2, 4, 8

hs-wred-queue

Syntax	hs-wred-queue [policy <i>slope-policy-name</i>] no hs-wred-queue
Context	config>qos>sap-egress>queue
Description	This command reverts the slope policy applied to the HSQ queue group queue to the default policy. Specifying an existing slope policy applies the named slope policy to the queue. The no form of the command reverts to the default slope policy.

Default	hs-wred-queue policy "_tmnx_hs_default"
Parameters	<i>slope-policy-name</i> — Specifies an existing slope policy to apply to this HSQ queue group queue.

hs-wrr-weight

Syntax	hs-wrr-weight <i>weight</i> no hs-wrr-weight
Context	config>qos>sap-egress>queue
Description	This command specifies the WRR relative weight, with which this queue should parent into an HSQ WRR group defined within the associated HS attachment policy. The weight of each queue determines how much bandwidth that queue gets out of the total rate for the HSQ WRR group. The no form of the command reverts to the default value.
Default	hs-wrr-weight 1
Parameters	<i>weight</i> — Specifies the HS WRR group queue weight. Values 1 to 127

15.4.2.8 Queue Group Template Policy HS Commands

hs-attachment-policy

Syntax	hs-attachment-policy <i>policy-name</i> no hs-attachment-policy
Context	config>qos>qgrps>egr>qgrp
Description	This command associates an existing HS attachment policy with the egress queue group template. The HS attachment policy controls how the egress queue group instance queues are attached to scheduler classes or WRR groups, and how WRR groups are attached to the scheduler classes. It also defines the mapping of the scheduling classes to the queues' aggregate shaper's low and high burst limit thresholds. Only one HS attachment policy can be associated with an egress queue group template. The no form of the command removes the policy name from the configuration and reapplies the default HS attachment policy.

Parameters *policy-name* — Specifies an existing attachment policy, up to 32 characters. Each HS attachment policy must be uniquely named within the system.

hs-wrr-group

Syntax **[no] hs-wrr-group** *group-id*

Context config>qos>qgrps>egr>qgrp

Description This command enables the context to configure HS WRR group information in the egress queue group template. The **hs-wrr-group** command is used to provision the rate and class weight of each of the two WRR scheduling groups that can be utilized by the egress queue group instance HSQ queues.

The **no** form of the command resets the HS WRR group parameters to their default values.

Parameters *group-id* — Specifies the HS WRR group identifier. WRR group ID 1 or 2 must be specified when executing the **hs-wrr-group** command. The specified group ID identifies which WRR group context is entered for editing.

Values 1, 2

adaptation-rule

Syntax **adaptation-rule** [*pir adaptation-rule*]
no adaptation-rule

Context config>qos>qgrps>egr>qgrp>hs-wrr-group

Description This command specifies how the system should resolve differences between the specified scheduling limit derived from the WRR group's rate command and the actual operational rate obtainable in hardware. The mutually exclusive **min**, **max**, and **closest** keywords specify whether the next highest rate, next lowest rate, or closest rate should be selected by the system.

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

Default adaptation-rule pir closest

Parameters *adaptation-rule* — Specifies the adaptation rule (**min**, **max**, or **closest**) to be used while computing the operational PIR value. The adaptation rule specifies the rules to compute the operational values while maintaining minimum offset. The **min**, **max**, and **closest** keywords are mutually exclusive.

Values min — When **min** is specified, the queue's rate parameter is treated as the minimum rate to shape the queue. The hardware chooses the appropriate timers and PIR leaky bucket behavior to ensure that the queues shaping rate is the closest possible value without going under the specified rate.

max — When **max** is specified, the queue's rate parameter is treated as the maximum rate to shape the queue. The hardware chooses the appropriate timers and PIR leaky bucket behavior to ensure that the queue's shaping rate is the closest possible value without going over the specified rate.

closest — When **closest** is specified, the queue's rate parameter is treated as the target rate to shape the queue. The hardware chooses the appropriate timers and PIR leaky bucket behavior to ensure that the queue's shaping rate is the closest possible value and can be higher or lower than the specified rate.

hs-class-weight

Syntax	hs-class-weight <i>weight</i> no hs-class-weight
Context	config>qos>qgrps>egr>qgrp>hs-wrr-group
Description	This command specifies the class weight of this WRR group at its parent primary shaper, relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class. This allows the capacity available at the primary shaper scheduling class to be shared in a WRR manner between the HSQ queue group queues and WRR groups attached to that scheduling class. The hs-class-weight parameter can be used to give unequal shares of the available capacity to different types of service offerings. The no form of the command reverts to the default value.
Default	hs-class-weight 1
Parameters	<i>weight</i> — Specifies the class weight of the HS WRR group. Values 1, 2, 4, 8

percent-rate

Syntax	percent-rate <i>percent</i> no percent-rate
Context	config>qos>qgrps>egr>qgrp>hs-wrr-group
Description	This command specifies the scheduling rate applied to the HS WRR group as a percentage of the port rate, including both the port's egress-rate and port's HS scheduler policy max-rate , if configured. The percent-rate and rate commands are mutually exclusive. The no form of the command reverts to the rate max.

Parameters *percent* — Specifies the percent rate of the HS WRR group.
Values 0.01 to 100.00

rate

Syntax **rate** *rate*
no rate

Context config>qos>qgrps>egr>qgrp>hs-wrr-group

Description This command specifies the scheduling rate applied to the HS WRR group in Kb/s. Alternatively, the keyword **max** can be specified, which removes the bandwidth limitation on the HS WRR group. The **rate** and **percent-rate** commands are mutually exclusive.

The **no** form of the command reverts to the **rate max**.

Default rate max

Parameters *rate* — Specifies the scheduling rates of the HS WRR group in Kb/s.
Values 1 to 2000000000, max

hs-alt-port-class-pool

Syntax [**no**] **hs-alt-port-class-pool**

Context config>qos>qgrps>egr>qgrp>queue

Description This command specifies that the HSQ queue group queues use the class buffers from the HS alternate port class buffer pools.

The **no** form of the command reverts to the HSQ queue group queues using buffers from HS standard port class pools.

hs-class-weight

Syntax **hs-class-weight** *weight*
no hs-class-weight

Context config>qos>qgrps>egr>qgrp>queue

Description This command specifies the class weight of this queue at its parent primary shaper, relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class. This allows the capacity available at the primary shaper scheduling class to be shared in a WRR manner between the HSQ queue group queues and WRR groups attached to that scheduling class. The **hs-class-weight** parameter can be used to give unequal shares of the available capacity to different types of service offerings.

This command is ignored for egress HSQ queue group queues, which are attached to an HS WRR group within an associated HS attachment policy. In this case the configuration of the **hs-class-weight** is performed under the **hs-wrr-group** within the egress queue group template.

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

Default	hs-class-weight 1
Parameters	<i>weight</i> — Specifies the class weight of the queue.
Values	1, 2, 4, 8

hs-wred-queue

Syntax	hs-wred-queue [<i>policy slope-policy-name</i>] no hs-wred-queue
Context	config>qos>qgrps>egr>qgrp>queue
Description	This command reverts the slope policy applied to the HSQ queue group queue to the default policy. Specifying an existing slope policy applies the named slope policy to the queue. The no form of the command reverts to the default slope policy.
Default	hs-wred-queue policy "_tmnx_hs_default"
Parameters	<i>slope-policy-name</i> — Specifies an existing slope policy name to apply to this HSQ queue group queue.

hs-wrr-weight

Syntax	hs-wrr-weight <i>weight</i> no hs-wrr-weight
Context	config>qos>qgrps>egr>qgrp>queue
Description	This command specifies the WRR relative weight, with which this queue should parent into an HSQ WRR group defined within the associated HS attachment policy. The weight of each queue determines how much bandwidth that queue gets out of the total rate for the HSQ WRR group. The no form of the command reverts to the default value.
Default	hs-wrr-weight 1
Parameters	<i>weight</i> — Specifies HS WRR group queue weight.
Values	1 to 127

15.4.2.9 Operational Commands

copy

Syntax	copy hs-attachment-policy <i>src-name dst-name</i> [overwrite] copy hs-pool-policy <i>src-name dst-name</i> [overwrite] copy hs-port-pool-policy <i>src-name dst-name</i> [overwrite] copy hs-scheduler-policy <i>src-name dst-name</i> [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	<i>src-name dst-name</i> — Indicates that the source policy ID and the destination policy ID are HS policy IDs. Specify the source policy ID that the copy command attempts to copy from and specify the destination policy ID to which the command copies a duplicate of the policy. overwrite — Specifies to replace the existing destination policy. Everything in the existing destination policy is overwritten with the contents of the source policy. If overwrite is not specified, an error occurs if the destination policy ID exists. <pre>SR>config>qos# copy hs-pool-policy policy1 policy2 MINOR: CLI Destination "policy2" exists use {overwrite}. SR>config>qos# copy hs-pool-policy policy1 policy2 overwrite</pre>

15.4.2.10 Show Commands

hs-pools

Syntax **hs-pools** *card-slot-number* **fp** *forwarding-plane* **egress**
hs-pools **port** *port-id* **egress**
hs-pools **port** *port-id* **egress** **network-queues**
hs-pools **port** *port-id* **egress** **queue-group** *queue-group-name* [**instance** *instance-id*]
hs-pools **port** *port-id* **egress** **sap** *sap-id*
hs-pools **port** *port-id* **egress** **subscriber** *sub-ident-string*

Context show

Description This command shows the egress pool information relating to an HSQ IOM.

The **card** and **fp** output shows the total buffer allocation, the number of allocated buffers, the available buffer allocation, and the buffer high watermarks for the system pools and the hierarchy of user-provisioned root and mid pools on the specified card and FP. The output includes the applied slope policy and the related instantaneous slope drop probabilities (as a percentage) for the user-provisioned pools.

The port output adds the port class pool information to the above hierarchy of user-provisioned pools for the specified port. The queue information for network queues, queue group instances, SAPs, or subscribers on the specified port is included when the associated parameter is added.

Parameters *card-slot-number* — Displays information for the specified card slot.

Values Depending on the chassis model, IOM slots are numbered from 1 to 10.

forwarding-plane — Displays information for the specified forwarding plane.

Values 1

port-id — Displays information about the specified physical port ID in the slot/mda/port format.

network-queues — Displays information about the egress network queues on the specified port.

queue-group-name — Displays information about the specified queue group name, up to 32 characters.

instance-id — Displays information about the specified queue group instance.

Values 1 to 65535

sap-id — Displays information about the specified SAP, up to 64 characters, on the specified port.

subscriber *sub-ident-string* — Displays information about the specified subscriber, up to 64 characters, on the specified port.

Output The following output is an example of HS pool information.

Sample Output

```
*A:PE4# show hs-pools port 1/1/2 egress sap 1/1/2:1

=====
HS Pools Port Information
=====
Port                : 1/1/2

-----
System Pool Information
-----
Total Buffers       : 209412 KB          Allocated           : 0 KB
Available           : 209412 KB          High Water Mark       : 0 KB
-----
Buffer Pool Hierarchy Information
-----

Root Pool : 1
| Total           : 2984148 KB  Allocated           : 0 KB
| Available        : 2984148 KB  High Water Mark     : 0 KB
| Hi-Slope Drop Prob : 0          Lo-Slope Drop Prob: 0
| Excd-Slope Drop Prob: 0
| Hs Slope Policy   : _tmnx_hs_default
|
| --- Mid Pool : 1
|   | Total           : 1193658 KB  Allocated           : 0 KB
|   | Available        : 1193658 KB  High Water Mark     : 0 KB
|   | Hi-Slope Drop Prob : 0          Lo-Slope Drop Prob: 0
|   | Excd-Slope Drop Prob: 0
|   | Hs Slope Policy   : _tmnx_hs_default
|   |
|   | --- Std Port Class Pool : 6
|   |   | Total           : 596828 KB  Allocated           : 0 KB
|   |   | Available        : 596828 KB  High Water Mark     : 0 KB
|   |   | Hi-Slope Drop Prob : 0          Lo-Slope Drop Prob: 0
|   |   | Excd-Slope Drop Prob: 0
|   |   | Hs Slope Policy   : _tmnx_hs_default
|   |   |
|   | --- Mid Pool : 2
|   |   | Total           : 1044450 KB  Allocated           : 0 KB
|   |   | Available        : 1044450 KB  High Water Mark     : 0 KB
|   |   | Hi-Slope Drop Prob : 0          Lo-Slope Drop Prob: 0
|   |   | Excd-Slope Drop Prob: 0
|   |   | Hs Slope Policy   : _tmnx_hs_default
|   |   |
|   |   | --- Std Port Class Pool : 5
|   |   |   | Total           : 522224 KB  Allocated           : 0 KB
|   |   |   | Available        : 522224 KB  High Water Mark     : 0 KB
|   |   |   | Hi-Slope Drop Prob : 0          Lo-Slope Drop Prob: 0
|   |   |   | Excd-Slope Drop Prob: 0
|   |   |   | Hs Slope Policy   : _tmnx_hs_default
|   |   |   |
```

```

--- Mid Pool : 3
| Total          : 895244 KB   Allocated          : 0 KB
| Available      : 895244 KB   High Water Mark   : 0 KB
| Hi-Slope Drop Prob : 0       Lo-Slope Drop Prob: 0
| Excd-Slope Drop Prob: 0
| Hs Slope Policy  : _tmnx_hs_default

--- Std Port Class Pool : 4
| Total          : 447622 KB   Allocated          : 0 KB
| Available      : 447622 KB   High Water Mark   : 0 KB
| Hi-Slope Drop Prob : 0       Lo-Slope Drop Prob: 0
| Excd-Slope Drop Prob: 0
| Hs Slope Policy  : _tmnx_hs_default

--- Mid Pool : 4
| Total          : 746036 KB   Allocated          : 0 KB
| Available      : 746036 KB   High Water Mark   : 0 KB
| Hi-Slope Drop Prob : 0       Lo-Slope Drop Prob: 0
| Excd-Slope Drop Prob: 0
| Hs Slope Policy  : _tmnx_hs_default

--- Std Port Class Pool : 1
| Total          : 373018 KB   Allocated          : 0 KB
| Available      : 373018 KB   High Water Mark   : 0 KB
| Hi-Slope Drop Prob : 0       Lo-Slope Drop Prob: 0
| Excd-Slope Drop Prob: 0
| Hs Slope Policy  : hs-slope-1

Root Pool : 2
| Total          : 994716 KB   Allocated          : 0 KB
| Available      : 994716 KB   High Water Mark   : 0 KB
| Hi-Slope Drop Prob : 0       Lo-Slope Drop Prob: 0
| Excd-Slope Drop Prob: 0
| Hs Slope Policy  : _tmnx_hs_default

--- Mid Pool : 5
| Total          : 795772 KB   Allocated          : 0 KB
| Available      : 795772 KB   High Water Mark   : 0 KB
| Hi-Slope Drop Prob : 0       Lo-Slope Drop Prob: 0
| Excd-Slope Drop Prob: 0
| Hs Slope Policy  : _tmnx_hs_default

--- Mid Pool : 6
| Total          : 198942 KB   Allocated          : 0 KB
| Available      : 198942 KB   High Water Mark   : 0 KB
| Hi-Slope Drop Prob : 0       Lo-Slope Drop Prob: 0
| Excd-Slope Drop Prob: 0
| Hs Slope Policy  : _tmnx_hs_default

```

Queue Information

```

Queue Name      : 1->1/1/2:1->1
FC Map          : be 12 11 h2 nc
Admin PIR       : 40000          Oper PIR           : 0
Admin MBS       : 64 KB          Oper MBS           : 64 KB
HS Wrr Group    : 1
HS Wrr Class Weight: 1          HS Wrr Weight      : 2

```

```

Depth                : 0
HS Class              : 1                      HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default
Queue Name           : 1->1/1/2:1->2
FC Map               : af
Admin PIR             : 40000                  Oper PIR                : 0
Admin MBS            : 64 KB                  Oper MBS                : 64 KB
HS Wrr Group         : 1
HS Wrr Class Weight  : 1                      HS Wrr Weight           : 1
Depth               : 0
HS Class             : 1                      HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default
Queue Name           : 1->1/1/2:1->6
FC Map               : ef
Admin PIR            : 20000                  Oper PIR                : 20000
Admin MBS            : 64 KB                  Oper MBS                : 64 KB
HS Wrr Group         : (not-applicable)
HS Wrr Class Weight  : 1                      HS Wrr Weight           : 0
Depth               : 0
HS Class             : 4                      HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default
Queue Name           : 1->1/1/2:1->7
FC Map               : h1
Admin PIR            : 10000                  Oper PIR                : 10000
Admin MBS            : 64 KB                  Oper MBS                : 64 KB
HS Wrr Group         : (not-applicable)
HS Wrr Class Weight  : 1                      HS Wrr Weight           : 0
Depth               : 0
HS Class             : 5                      HS Alt Port Class Pool : No
HS Slope Policy       : tmnx hs default

```

port

Syntax	port <i>port-id</i> hs-secondary-shaper port <i>port-id</i> hs-secondary-shaper <i>secondary-shaper-name</i> [statistics associations]
Context	show
Description	This command displays the details of the HS secondary shapers configured on the specified port.
Parameters	<i>port-id</i> — Displays information about the specified physical port ID in the slot/mda/port format. <i>secondary-shaper-name</i> — Displays information about the HS secondary name, up to 32 characters. statistics — Displays statistics for the HS secondary shaper. The <i>secondary-shaper-name</i> must be specified. associations — Displays the entities associated with the HS secondary shaper. The <i>secondary-shaper-name</i> must be specified.

Output The following output is an example of port HS secondary shaper information.

Sample Output

```
*A:PE# show port 1/1/2 hs-secondary-shaper "default"
```

```
=====
Ethernet Port 1/1/2 Egress HS Secondary Shaper Information
=====
Policy Name      : default
Description      : (Not Specified)
Rate             : max
Low Burst Max Class: 6
```

```
-----
Class            Rate
-----
1                max
2                max
3                max
4                max
5                max
6                max
-----
```

```
*A:PE# show port 1/1/2 hs-secondary-shaper "default" statistics
```

```
=====
Ethernet Port 1/1/2 Egress HS Secondary Shaper Information
=====
Policy Name      : default
```

```
-----
Statistics Information
-----
```

```
-----
Packets          Octets
-----
Class 1
  Forwarded      : 0          0
Class 2
  Forwarded      : 0          0
Class 3
  Forwarded      : 0          0
Class 4
  Forwarded      : 0          0
Class 5
  Forwarded      : 0          0
Class 6
  Forwarded      : 0          0
Aggregate
  Forwarded      : 0          0
-----
```

```
*A:PE# show port 1/1/2 hs-secondary-shaper "default" associations
```

```

=====
Ethernet Port 1/1/2 Egress HS Secondary Shaper Information
=====
Policy Name      : default
Description      : (Not Specified)
Rate             : max
Low Burst Max Class: 6

-----
Class            Rate
-----
1                max
2                max
3                max
4                max
5                max
6                max
-----

-----
Service Associations
-----
Service ID      Service Type      SAP
-----
2                IES                1/1/2:2
-----

-----
Subscriber Associations
-----
Subscriber ID
-----
No Subscriber Associations Found.
-----
=====
*A:PE#

```

hs-attachment-policy

- Syntax** **hs-attachment-policy** [*policy-name*] [**association** | **detail**]
- Context** show>qos
- Description** This command displays information about HS attachment policies in the system.
- Parameters** *policy-name* — Displays information about the specified policy name, up to 32 characters.
- association** — Displays the entities associated with all HS attachment policies or the specified HS attachment policy.
- detail** — Displays detailed policy information, including policy associations, of all HS attachment policies, or the specified HS attachment policy.
- Output** The following output is an example of HS attachment policy information.

Sample Output

```
*A:PE# show qos hs-attachment-policy "default"
```

```
=====
HS Attachment Policy Information
=====
Policy Name           : default
Description            : Default hs attachment QoS policy
Low Burst Max Class    : 6

-----
Queue                 Scheduling Class          WRR Group
-----
1                     (Not-Applicable)           1
2                     (Not-Applicable)           1
3                     (Not-Applicable)           1
4                     2                      (Not-Applicable)
5                     3                      (Not-Applicable)
6                     4                      (Not-Applicable)
7                     5                      (Not-Applicable)
8                     6                      (Not-Applicable)

-----
WRR Group             Scheduling Class
-----
1                     1
2                     unattached
=====
```

```
*A:PE# show qos hs-attachment-policy "hs-att-pol-1" detail
```

```
=====
HS Attachment Policy Information
=====
Policy Name           : hs-att-pol-1
Description            : (Not Specified)
Low Burst Max Class    : 1

-----
Queue                 Scheduling Class          WRR Group
-----
1                     (Not-Applicable)           1
2                     (Not-Applicable)           1
3                     unattached                unattached
4                     unattached                unattached
5                     unattached                unattached
6                     4                      (Not-Applicable)
7                     5                      (Not-Applicable)
8                     6                      (Not-Applicable)

-----
WRR Group             Scheduling Class
-----
1                     1
2                     unattached
=====

Associations
```

```

-----
Network-Queue Policy
-----
10

Sap-Egress Policy
-----
10
20

Egress Queue-Group Templates
-----
queue-group-1

-----
=====
*A:PE#

```

hs-pool-policy

Syntax	hs-pool-policy [<i>policy-name</i>] [association detail]
Context	show>qos
Description	This command displays information about HS pool policies in the system.
Parameters	<i>policy-name</i> — Displays information about the specified HS pool policy name, up to 32 characters. association — Displays the entities associated with all HS pool policies or the specified HS pool policy. detail — Displays detailed policy information, including policy associations, of all HS pool policies or the specified HS pool policy.
Output	The following output is an example of HS pool policy information.

Sample Output

```

*A:PE# show qos hs-pool-policy

=====
HS Pool Policy Summary
=====
Policy Name          Description
-----
default              Default hs pool QoS policy
hs-pool-pol-1        (Not Specified)
-----
No. of Policies: 2
=====

*A:PE# show qos hs-pool-policy default detail

=====

```

HS Pool Policy Information

```
=====
Policy Name      : default
Description      : Default hs pool QoS policy
System Reserve   : 5.00
```

----- Root Pool Information

```
-----
Pool Id          : 1           Allocation Weight : 75
Slope Policy     : _tmnx_hs_default
Pool Id          : 2           Allocation Weight : 25
Slope Policy     : _tmnx_hs_default
Pool Id          : 3           Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 4           Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 5           Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 6           Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 7           Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 8           Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 9           Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 10          Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 11          Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 12          Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 13          Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 14          Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 15          Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
Pool Id          : 16          Allocation Weight : 0
Slope Policy     : _tmnx_hs_default
```

----- Mid Pool Information

```
-----
Pool Id          : 1           Allocation Percent : 40.00
Port BW Oversub Factor : 1       Parent Root Pool : 1
Slope Policy     : _tmnx_hs_default
Pool Id          : 2           Allocation Percent : 35.00
Port BW Oversub Factor : 1       Parent Root Pool : 1
Slope Policy     : _tmnx_hs_default
Pool Id          : 3           Allocation Percent : 30.00
Port BW Oversub Factor : 1       Parent Root Pool : 1
Slope Policy     : _tmnx_hs_default
Pool Id          : 4           Allocation Percent : 25.00
Port BW Oversub Factor : 1       Parent Root Pool : 1
Slope Policy     : _tmnx_hs_default
```



```

Pool Id           : 5           Allocation Percent : 80.00
Port BW Oversub Factor : 1           Parent Root Pool : 2
Slope Policy      : _tmnx_hs_default
Pool Id           : 6           Allocation Percent : 20.00
Port BW Oversub Factor : 1           Parent Root Pool : 2
Slope Policy      : _tmnx_hs_default
Pool Id           : 7           Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default
Pool Id           : 8           Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default
Pool Id           : 9           Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default
Pool Id           : 10          Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default
Pool Id           : 11          Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default
Pool Id           : 12          Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default
Pool Id           : 13          Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default
Pool Id           : 14          Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default
Pool Id           : 15          Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default
Pool Id           : 16          Allocation Percent : 1.00
Port BW Oversub Factor : 1           Parent Root Pool : 0
Slope Policy      : _tmnx_hs_default

```

```

-----
-----
Card Forwarding Plane (FP) Associations
-----
Card          FP
-----
1              1
-----
=====
*A:PE#

```

hs-port-pool-policy

Syntax **hs-port-pool-policy** [*policy-name*] [**association** | **detail**]

Context show>qos

- Description** This command displays information about HS port pool policies in the system.
- Parameters** *policy-name* — Displays information about the specified HS port pool policy name, up to 32 characters.
- association** — Displays the entities associated with all HS port pool policies or the specified HS port pool policy.
- detail** — Displays detailed policy information, including policy associations, of all HS port pool policies, or the specified HS port pool policy.
- Output** The following output is an example of HS port pool policy information.

Sample Output

```
*A:PE# show qos hs-port-pool-policy

=====
HS Port Pool Policy Summary
=====
Policy Name          Description
-----
default              Default hs port pool QoS policy
no-class-pools       (Not Specified)
hs-port-pool-pol-1   (Not Specified)
-----
No. of Policies: 3
=====
*A:PE# show qos hs-port-pool-policy default detail

=====
HS Port Pool Policy Information
=====
Policy Name          : default
Description           : Default hs port pool QoS policy
-----

Standard Port Class Pool Information
-----
Class Id             : 1          Parent Mid Pool      : 1
Alloc Port BW Weight : 1          Alloc Explicit Prcnt: 0.00
Slope Policy          : _tmnx_hs_default
Class Id             : 2          Parent Mid Pool      : 2
Alloc Port BW Weight : 1          Alloc Explicit Prcnt: 0.00
Slope Policy          : _tmnx_hs_default
Class Id             : 3          Parent Mid Pool      : 3
Alloc Port BW Weight : 1          Alloc Explicit Prcnt: 0.00
Slope Policy          : _tmnx_hs_default
Class Id             : 4          Parent Mid Pool      : 4
Alloc Port BW Weight : 1          Alloc Explicit Prcnt: 0.00
Slope Policy          : _tmnx_hs_default
Class Id             : 5          Parent Mid Pool      : 5
Alloc Port BW Weight : 1          Alloc Explicit Prcnt: 0.00
Slope Policy          : _tmnx_hs_default
Class Id             : 6          Parent Mid Pool      : 6
Alloc Port BW Weight : 1          Alloc Explicit Prcnt: 0.00
Slope Policy          : _tmnx_hs_default
```

```

-----
-----
Alternate Port Class Pool Information
-----
Class Id           : 1           Parent Mid Pool       : 0
Alloc Port BW Weight : 1           Alloc Explicit Prcnt: 0.00
Slope Policy       : _tmnx_hs_default
Class Id           : 2           Parent Mid Pool       : 0
Alloc Port BW Weight : 1           Alloc Explicit Prcnt: 0.00
Slope Policy       : _tmnx_hs_default
Class Id           : 3           Parent Mid Pool       : 0
Alloc Port BW Weight : 1           Alloc Explicit Prcnt: 0.00
Slope Policy       : _tmnx_hs_default
Class Id           : 4           Parent Mid Pool       : 0
Alloc Port BW Weight : 1           Alloc Explicit Prcnt: 0.00
Slope Policy       : _tmnx_hs_default
Class Id           : 5           Parent Mid Pool       : 0
Alloc Port BW Weight : 1           Alloc Explicit Prcnt: 0.00
Slope Policy       : _tmnx_hs_default
Class Id           : 6           Parent Mid Pool       : 0
Alloc Port BW Weight : 1           Alloc Explicit Prcnt: 0.00
Slope Policy       : _tmnx_hs_default
-----
-----
Port Ethernet Egress Associations
-----
1/2/1
-----
=====
*A:PE#

```

hs-scheduler-hierarchy

Syntax	hs-scheduler-hierarchy port <i>port-id</i> [hs-secondary-shaper <i>shaper-name</i>] [interval <i>time-in-seconds</i>] hs-scheduler-hierarchy port <i>port-id</i> [interval <i>time-in-seconds</i>] hs-secondary-shapers hs-scheduler-hierarchy port <i>port-id</i> [interval <i>time-in-seconds</i>] queue-group <i>queue-group-name</i> instance <i>instance-id</i> {access network} hs-scheduler-hierarchy sap <i>sap-id</i> egress [interval <i>time-in-seconds</i>] hs-scheduler-hierarchy subscriber <i>sub-ident</i> egress [interval <i>time-in-seconds</i>]
Context	show>qos
Description	This command displays the egress HS scheduler hierarchy for the HS scheduler policy applied to the related port. The output includes the current maximum rate and each scheduling class rate on the port. The parameters allow for the current rates for the scheduling classes for all HS secondary shapers or a specific HS secondary shaper, or for the queue rates in a specified access or network queue group instance, SAP or subscriber to be included in the output.

- Parameters**
- port-id* — Specifies the physical port ID in the slot/mda/port format.
 - shaper-name* — Displays information about the specified HS secondary name, up to 32 characters, and includes the scheduling class rates for the HS secondary shaper on the port.
 - time-in-seconds* — Displays the time interval used to calculate the rates, in seconds.
 - Values** 1 to 5
 - hs-secondary-shapers** — Displays information about scheduling class rates for all HS secondary shapers on the port.
 - queue-group-name* — Displays information about the specified queue group name on the port, up to 32 characters.
 - instance-id* — Displays information about the specified queue group instance.
 - Values** 1 to 65535
 - sap-id* — Displays information about the queue rates for the specified SAP, up to 64 characters, on the specified port.
 - sub-ident* — Displays information about the specified subscriber, up to 64 characters, on the specified port.

Output The following output is an example of HS scheduler hierarchy information.

Sample Output

```
*A:PE# show qos hs-scheduler-hierarchy port 1/1/1 hs-secondary-shaper "default"

=====
Hs Scheduler Hierarchy Information
=====
Hs Sched Policy Name          : hs-sched-pol-1

Port Max-Rate : 0 Mbps
Hs-Sec-Shaper:default Agg-Rate : 0 Kbps

Scheduler Priority 6
  Scheduler Class 6 Rate : 0 Mbps
    Hs-Sec-Shaper:default Class 6 Rate : 0 Kbps

Scheduler Priority 5
  Scheduler Class 5 Rate : 0 Mbps
    Hs-Sec-Shaper:default Class 5 Rate : 0 Kbps

Scheduler Priority 4
  Scheduler Class 4 Rate : 0 Mbps
    Hs-Sec-Shaper:default Class 4 Rate : 0 Kbps

Scheduler Priority 3
  Scheduler Class 3 Rate : 0 Mbps
    Hs-Sec-Shaper:default Class 3 Rate : 0 Kbps

Scheduler Priority 2
  Scheduler Class 2 Rate : 0 Mbps
    Hs-Sec-Shaper:default Class 2 Rate : 0 Kbps
```

```
Scheduler Priority 1
  Scheduler Class 1 Rate : 0 Mbps
    Hs-Sec-Shaper:default Class 1 Rate : 0 Kbps
=====
*A:PE# show qos hs-scheduler-hierarchy port 1/1/1 hs-secondary-shapers

=====
Hs Scheduler Hierarchy Information
=====
Hs Sched Policy Name      : hs-sched-pol-1

Port Max-Rate : 0 Mbps

Scheduler Priority 6
  Scheduler Class 6 Rate : 0 Mbps

Scheduler Priority 5
  Scheduler Class 5 Rate : 0 Mbps

Scheduler Priority 4
  Scheduler Class 4 Rate : 0 Mbps

Scheduler Priority 3
  Scheduler Class 3 Rate : 0 Mbps

Scheduler Priority 2
  Scheduler Class 2 Rate : 0 Mbps

Scheduler Priority 1
  Scheduler Class 1 Rate : 0 Mbps

-----
HS Secondary Shaper Rates
-----
Hs-Sec-Shaper:default Agg-Rate : 0 Kbps
  Class 6 Rate : 0 Kbps
  Class 5 Rate : 0 Kbps
  Class 4 Rate : 0 Kbps
  Class 3 Rate : 0 Kbps
  Class 2 Rate : 0 Kbps
  Class 1 Rate : 0 Kbps

-----
=====
*A:PE#
```

hs-scheduler-policy

Syntax	hs-scheduler-policy [<i>policy-name</i>] [association detail]
Context	show>qos
Description	This command displays information about HS scheduler policies in the system.

- Parameters**
- policy-name* — Displays information about the specified HS scheduler policy name, up to 32 characters.
 - association** — Displays the entities associated with all HS scheduler policies or the specified HS pool policy.
 - detail** — Displays the policy information, including policy associations, of all HS scheduler policies or the specified HS scheduler policy.
- Output** The following output is an example of HS scheduler policy information.

Sample Output

```
*A:PE# show qos hs-scheduler-policy

=====
HS Scheduler Policy Summary
=====
Policy Name          Description
-----
default              Default hs scheduler QoS policy
hs-sched-pol-1       (Not Specified)
-----
No. of Policies: 2
=====

*A:PE# show qos hs-scheduler-policy "default" detail

=====
HS Scheduler Policy Information
=====
Policy Name          : default
Description           : Default hs scheduler QoS policy
Max Rate              : max

-----
Scheduling Class    Rate          Group          Weight in Group
-----
1                   max            0                1
2                   max            0                1
3                   max            0                1
4                   max            0                1
5                   max            0                1
6                   max            0                1

-----
Group              Rate
-----
1                  max

-----
Port Ethernet Egress Associations

-----
1/1/2
1/1/3
1/1/4
1/1/5
```

1/1/6
1/1/7
1/1/8
1/1/9
1/1/10
1/2/1

*A: PE#

15.4.2.11 Clear Commands

card

Syntax	card slot-number fp fp-number hs-pool high-water-mark card slot-number fp fp-number hs-pool high-water-mark mid-pool [value] card slot-number fp fp-number hs-pool high-water-mark root-pool [value] card slot-number fp fp-number hs-pool high-water-mark system
Context	clear
Description	This command re-initializes the card in the specified slot. A clear card command (without the soft) is referred to as a <i>Hard Reset</i> . A clear card x soft command (with the soft) is referred to as a <i>Soft Reset</i> .
Parameters	<i>slot-number</i> — Clears information for the specified card slot. Values 1 to 20, A, B, C, D <i>fp-number</i> — Displays information for the specified forwarding plane. Values 1 to 6 high-water-mark — Clears the high water marks within the HS pool on the specified card and FP. mid-pool — Clears the high water marks within the mid pools in the HS pool and on the specified card and FP. Values 1 to 16 root-pool — Clears the high water marks within the root pools in the HS pool and on the specified card and FP. Values 1 to 16 system — Clears the high water marks within the in the HS system pool and on the specified card and FP.

port

Syntax	port <i>port-id</i> hs-secondary-shaper <i>shaper-name</i> statistics port <i>port-id</i> hs-pool high-water-mark {[standard value] [alternate value]}
Context	clear
Description	This command clears port statistics for the specified ports.
Parameters	<i>port-id</i> — Specifies the physical port identifier. Values <i>slot</i> [/<i>mda</i>[/<i>port</i>]] or <i>slot/mda/port</i> [<i>channel</i>]] <i>shaper-name</i> — Clears information about the specified shaper name, up to 32 characters. standard value — Clears high water marks within specific standard port class pools using the standard parameters. Values 1 to 6 alternate value — Clears high water marks within specific alternative port class pools using the alternative parameters. Values 1 to 6

16 QoS in MC-MLPPP

16.1 Overview

If the user enables the multi-class option under an MLPPP bundle, the XMA or MDA egress data path provides a queue for each of the four 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, therefore, in which of the four egress XMA or 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 described in [Table 121](#).

Table 121 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 122](#) describes a different mapping enabled when the user applies one of three predefined egress QoS profiles in the 4-class bundle configuration only.

Table 122 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

Table 122 Packet Forwarding Class to MLPPP Class Mapping (Continued)

FC ID	FC Name	Scheduling Priority (Default)	MLPPP Class (MLPPP Egress QoS profile 1, 2, and 3)
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 predefined egress QoS profiles to an MLPPP bundle.

[Table 123](#) and [Figure 52](#) describe 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 123 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

Table 123 MLPPP Class Queue Threshold Parameters (Continued)

	Class 0		Class 1		Class 2		Class 3	
4-Class Bundle Egress QoS Profile 3	25	12	5	3	200	100	1000	500

Figure 52 MLPPP Class Queue Thresholds for In-Profile and Out-of-Profile Packets





Table 124 and Figure 53 describe the class queue scheduling parameters.

Table 124 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 53 MLPPP Class Queue Scheduling Scheme



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 2 and 3 queues 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.

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 serviced with strict priority in ascending order of class number.

16.2 Ingress MLPPP Class Reassembly

For an MLPPP bundle with the multi-class option enabled, there is a default profile for setting the reassembly timer value for each class. When the predefined MLPPP ingress QoS profile 1 is applied to a 4-class bundle, the values of the timers are modified as described in [Table 125](#).

Table 125 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

16.3 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

```

16.3.1 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

```

16.3.2 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 describes the parameters and the configuration processes and rules.

- The user creates an ingress QoS profile, mlppp-profile-ingress, to configure the desired value of the ingress per-class reassembly timer. Ingress QoS profile #1 is reserved to the predefined profile with parameter values in [Table 125](#). 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 125](#) for the ingress QoS Profile #1, regardless what parameter value the edited Profile #1 has at that time.

- 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 predefined profiles with parameter values shown in [Table 122](#), [Table 123](#), or [Table 124](#). 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 122](#), [Table 123](#), or [Table 124](#) for the egress QoS Profile #1. This is regardless of the parameter value the edited profiles have at that time.
- A maximum of 128 ingress QoS profiles and 128 egress QoS profiles can be created on the system.
- The values of the ingress per-class reassembly timer are configured in the ingress QoS profile.
- 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 122](#) for the mapping when one of the three predefined 4-class egress QoS profiles is selected.
- 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 52](#) and as `max-queue-size` in CLI. The out-of-profile threshold for an MLPPP class queue, referred to as `oop` in [Figure 52](#), is not directly configurable and is set to 50% of the maximum queue size rounded up to the nearest higher integer value.
- The MLPPP class queue scheduling parameters are configured in the egress QoS profile. The minimum information rate, referred to as `MIR` in [Figure 53](#) 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 53](#) and `weight` in CLI, applies to class 1, class 2, and class 3 queues. `W1` in [Figure 53](#) 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 as percentage Class 2 queue shares and Class 3 queue shares of the available bundle rate.
- 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.
- 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.

- The following operations can be performed without shutting down the associated bundles:
 - Changes to any parameters in the ingress and egress QoS profiles.

16.4 QoS in MLFR and FRF.12 Fragmentation

The following sections describe MLFR and FRF.12 feature descriptions and implementation.

16.4.1 QoS in MLFR

The MLFR feature introduces the following new QoS requirements on the XM or MDA:

- Four XMA or 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 XMA or 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 126](#) and [Figure 54](#) describe 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 126 Default FR Class Queue Threshold Parameters

	Class 0		Class 1		Class 2		Class 3	
	Max	Oop	Max	Oop	Max	Oop	Max	Oop
Queue threshold (in ms@available bundle rate)	10	5	50	25	150	75	750	375

Figure 54 FR Class Queue Thresholds for In-Profile and Out-of-Profile Packets





Table 127 and Figure 55 describe the class queue scheduling parameters for an MLFR bundle.

Table 127 Default FR Class Queue Scheduling Parameters

WRR Parameters			
MIR	W1	W2	W3
90%	<1%	89%	10%

Figure 55 FR Class Queue Scheduling for an MLFR Bundle



The minimum information rate, referred to as MIR in [Figure 55](#) 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 127](#) and weight in CLI, applies to class 1, class 2, and class 3 queues. 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.

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, the user can configure the value of the FR scheduling class ingress reassembly timeout for an MLFR bundle. The default values of the timers are shown in [Table 128](#).

Table 128 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.

16.4.2 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 XMA or MDA is reusing the existing FR SAP high- and low-priority queues, as shown in [Figure 56](#).

Figure 56 DLC Egress Channel Queue Scheduling



al_0196

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 XMA or MDA queues should be used to queue the fragments. Classes 0 and 1 map all fragments of the FR SAP to the high-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: $SAP\_HP\_Q\_PIR = \text{Sum}\{\text{all } SAP\_FC\_Q\_PIR\}$
- Low priority (LP) SAP queue: $SAP\_LP\_Q\_WRR\_Weight = \text{Sum}\{\text{all } SAP\_FC\_Q\_CIR\}$

16.5 MLPPP Command Reference

16.5.1 Command Hierarchies

16.5.1.1 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
```

16.5.2 Command Descriptions

16.5.2.1 Configuration Commands

16.5.2.1.1 MC-MLPPP Commands

mlppp-profile-egress

Syntax	[no] mlppp-profile-egress <i>profile-id</i> [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	<i>profile-id</i> — Specifies an ingress mlppp-profile ID.
Values	1 to 65535

class

Syntax	class <i>class-id</i>
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	<i>class-id</i> — Specifies a class for this policy.
Values	0 to 3

max-queue-size

Syntax	max-queue-size <i>queue-size</i> 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	<i>queue-size</i> — Specifies the maximum queue size.
Values	1 to 1000 in ms of buffer space

mir

Syntax	mir <i>mir-value</i> 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	<i>mir-value</i> — Specifies the MIR scheduling parameter.
Values	1 to 100

weight

Syntax	weight <i>weight-value</i> 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	<i>weight-value</i> — Specifies the weight scheduling parameter.
Values	1 to 100

description

Syntax	description <i>description-string</i> 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, 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.

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 to 3

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 multiclass 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 to 65535

reassemble-timeout

Syntax	reassemble-timeout <i>timeout-value</i> no reassemble-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	<i>timeout-value</i> — Specifies a reassembly timeout for this policy. Values 1 to 1000 ms

17 Class Fair Hierarchical Policing (CFHP)

17.1 Overview

CFHP merges the benefits of non-delay rate enforcement inherent to policers with the priority and fairness sensitivity of queuing and scheduling. CFHP is implemented as a group of child policers mapped to a parent policer where the rate enforced by the parent both obeys strict priority levels and is class fair within a priority level. At the parent policer, the output of a lower priority child policer cannot prevent forwarding of packets of a higher priority child policer and when multiple child policers share the same priority level, the system maintains a Fair Information Rate (FIR) for each child that is separate from a child's PIR and CIR rates. Policers can also be used standalone. The parent is optional.

Multiservice sites support policer-control-policy in the in the ingress and egress in addition to scheduler-policy.

Below are the capabilities and limitations for CFHP under a multiservice site:

- Support for SAP only (no subscriber support)
- Assignment is for port only (not for card)
- Supported both in Ingress and Egress
- Policer Overrides are not supported under a multiservice site.

```
*A:Dut-A>config>service>cust>multi-service-site# pwc
```

```
-----  
Present Working Context :  
-----
```

```
<root>  
configure  
service  
customer 2  
multi-service-site "mss1"  
-----
```

```
*A:Dut-A>config>service>cust>multi-service-site# info
```

```
-----  
assignment port 9/1/4  
ingress  
policer-control-policy "pcp"  
exit  
egress  
policer-control-policy "pcp"  
exit  
-----
```

Example of a service using mss is as follows. The sap-egress qos policy “3” will have policers parented to arbiters that are configured in the policer-control-policy “pcp” as in the preceding example.

```
*A:Dut-A>config>service>vpls# pwc
```

```
-----  
Present Working Context :  
-----
```

```
<root>  
configure  
service  
vpls "101"
```



```
-----  
*A:Dut-A>config>service>vpls# info
```

```
-----  
shutdown  
stp  
shutdown  
exit  
sap 9/1/4 create  
multi-service-site "mss1"  
egress  
qos 3  
exit  
exit  
-----
```

17.2 Parent Policer Priority and Unfair Sensitive Discard Thresholds

Priority-level bandwidth control is managed on the parent policer through the use of progressively higher discard thresholds for each in-use priority level. Up to eight priority levels are supported and are individually enabled per parent policer instance based on child policer priority level association. When multiple child policers are associated with a parent policer priority level, two separate discard thresholds are maintained for that priority level. A lower “discard-unfair” threshold ensures that when a child policer has exceeded its FIR rate, its unfair packets are discarded first (assuming the parent policer’s bucket depth has reached the priority level’s “discard-unfair” threshold) protecting the priority level’s fair traffic from the priority level’s unfair traffic.

A second “discard-all” threshold is used to discard all remaining packets associated with the priority level in the case where higher priority traffic exists and the sum of both the priority level’s traffic and the higher priority traffic exceeds the parent policer rate. This protects the higher priority traffic on the parent policer from being discarded due to lower priority traffic. The child and parent policers operate in an atomic fashion, any conformance effect on a child policer’s bucket depth is canceled when the parent policer discards a packet. Policer bucket rate and packet flow interaction with bucket depth are shown in [Figure 57](#). Parent policer bucket and priority thresholds are shown in [Figure 58](#).

Figure 57 Policer Bucket Rate and Packet Flow Interaction with Bucket Depth**Figure 58** Parent Policer Bucket and Priority Thresholds

17.3 CFHP Ingress and Egress Use Cases

While ingress CFHP seems a natural fit based on how policers are typically used in today's networks, CFHP may also be used at egress. The reasons for utilizing egress CFHP may be to provide a non-jitter or latency inducing aggregate SLA for multiple ingress flows or just to provide higher scale in the number of egress aggregate SLAs supported.

17.4 Post-CFHP Queuing and Scheduling

Although CFHP enforces aggregate rate limiting while maintaining sensitivity to strict priority and fair access to bandwidth within a priority, CFHP output packets still require queuing and scheduling to provide access to the switch fabric or to an egress port.

17.4.1 Ingress CFHP Queuing

At ingress, CFHP output traffic is automatically mapped to a unicast or multipoint queue in order to reach the proper switch fabric destinations. In order to manage this automatic queuing function, a shared queue policy has been created or exists by default and is named `policer-output-queues`. For modifying parameters in this shared-queue policy, refer to [Shared-Queue QoS Policy Command Reference](#).

The unicast queues in the policy are automatically created on each destination switch fabric tap and ingress CFHP unicast packets automatically map to one of the queues based on forwarding class and destination tap. The multipoint queues within the policy are created on the IOM3-XP's 16/IMM or XMA multicast paths; 16 multicast paths are supported by default with 28 on 7950 XRS systems and the 7750 SR 12-e systems, with the latter having setting "tools perform the system set-fabric-speed fabric-speed-b." The multicast paths represent an available multicast switch fabric path; the number of each being controlled using the command:

CLI Syntax:

```
configure mcast-management bandwidth-policy policy-name
    t2-paths secondary-path
    number-paths number-of-paths [dual-sfm number-of-paths]
```

For ingress CFHP multicast packets (Broadcast, Unknown unicast, or Multicast—referred to as BUM traffic), the system maintains a conversation hash table per forwarding class and populates the table's forwarding class hash result entry with the one of the multicast paths. Best-effort traffic uses the secondary paths, and expedited traffic uses the primary paths. When a BUM packet is output by ingress CFHP, a conversation hash is performed and used along with the packet's forwarding class to select a hash table entry in order to derive the multicast path to be used. Each table entry maintains a bandwidth counter that is used to monitor the aggregate traffic per multicast path. The process can be optimized by enabling IMPM on any forwarding complex, which allows the system to redistribute this traffic across the IMPM paths on all forwarding complexes to achieve a more even capacity distribution. Be aware that enabling IMPM causes routed and VPLS (IGMP and PIM) snooped IPv4 multicast groups, and routed and PIM snooped (with **sg-based** forwarding) IPv6 multicast groups to be managed by IMPM.

Any discards performed in the ingress shared queues will be reflected in the ingress child policer's discard counters and reported statistics, assuming a discard counter capable stat-mode is configured for the child policer.

17.4.2 Egress CFHP Queuing

When CFHP is being performed at egress, queuing of the CFHP output packets is accomplished through egress queue group queues. The system maintains a special egress queue group template (policer-output-queues) that is automatically applied to all Ethernet access ports that are up. The number of queues, queue types (expedite or best-effort), queue parameters, and the default forwarding class mappings to the queues are managed by the template. On each Ethernet port, the queue parameters may be overridden.

When a SAP egress QoS policy is applied to an Ethernet SAP and the policy contains a forwarding class mapping to a CFHP child policer, the default behavior for queuing the CFHP output is to use the egress Ethernet port's policer-output-queues queue group and the forwarding class mapping within the group to choose the egress queue. Optionally, the SAP egress QoS policy may also explicitly define which egress queue to use within the default queue group or even map the policer output to a different, explicitly created queue group on the port.

Any discards performed in the egress queue group queues will be reflected in the egress child policer's discard counters and reported statistics assuming a discard counter capable stat-mode is configured for the child policer. Exceed-profile traffic from the policer will be counted as out-of-profile traffic in the egress queue group queues.

17.4.2.1 Policer to Local Queue Mapping

Egress policers can be optionally mapped to a local queue instead of a queue group queue where required.

The syntax for assigning one such egress policer mapped to local queue is as below:

```
*A:Dut-A>config>qos>sap-egress$ pwc
```

```
-----  
Present Working Context :
```

```
-----  
<root>  
configure  
qos  
sap-egress 3 create  
-----
```

```
*A:Dut-A>config>qos>sap-egress$ info
-----
queue 1 create
exit
queue 2 create
exit
policer 2 create
exit
fc ef create
policer 2 queue 2
exit
-----
```

To a local queue, as in "queue 2" in the previous example, both a policer and also a forwarding class can be concurrently mapped as shown below:

```
*A:Dut-A>config>qos>sap-egress$ info
-----
queue 1 create
exit
queue 2 create
exit
policer 2 create
exit
fc af create
queue 2
exit
fc ef create
policer 2 queue 2
exit
-----
```

A queue resource is allocated whenever there is either an fc or a policer referencing it. The local queue is freed when there are no references to it. The local queue cannot be deleted when it is being referenced. Exceed-profile traffic from the policer will be counted as out-of-profile traffic in the egress local queues.

17.4.3 Egress Subscriber CFHP Queuing

When a subscriber packet is mapped to a child policer through the SAP egress QoS policy, the actual egress queue group is derived from the subscriber host identification process within the subscriber management module; otherwise the default queue-group is used.

When egress policed packets are directed to a local SAP queue, and, when this is configured, the output of a **show service id id sap sap-id sap-stats** only counts these packets through the policer; they are not counted a second time through the queue to avoid double counting. Consequently, any packets sent directly (not through a policer) to a local SAP post-policer queue are not counted in the sap-stats output. The output of **show service id id sap sap-id stats** always counts these packets in both the related policer and queue. If it is required to count packets sent directly to the local SAP post-policer queue in the sap-stats output, the packets could be sent into a policer with the rate set to max, then into the local SAP queue.

17.4.3.1 Subscriber Destination String Queue Group Identification

When a subscriber is identified, a special destination string may optionally exist for the subscriber that is typically used to identify the subscriber's destination aggregation node. This feature applies only to the 7450 ESS and 7950 SR.

On the subscriber's egress Ethernet port, the default policer-output-queues and other explicitly created queue groups may be configured to represent a destination node by defining the same destination string on the queue group. When the subscriber's destination string is defined, the system will search the subscriber's egress port for an egress queue group with the same string defined. If found, it will use that matched queue group instead of the default queue group. If a queue-group matching the string is not found, the subscriber identification event will not fail and the subscriber host will be mapped to default policer-output-queues.

The destination node-based queuing model is designed to provide the ability to shape the aggregate subscriber output to a destination aggregation node based on a queue group created for the specific purpose. On the queue group, a scheduling-policy is applied that defines the desired virtual scheduling behavior of the queues and aggregate maximum rate of the queue group. The destination string matching function could be used to represent any arbitrary downstream bandwidth limit, not just an aggregation node. If the destination string is not present (null value), the default policer egress queue group ('policer-output-queues') on the subscriber's port will be used.

17.4.4 SAP Default Destination String

In order to simplify subscriber destination string provisioning, a **def-inter-dest-id** command can be used in a sub-sla-mgmt node within a SAP, which allows the definition of a default destination string for all subscribers associated with the SAP. The command also accepts the use-top-q flag that automatically derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation. This feature applies only to the 7450 ESS and 7750 SR.

The command is also supported within the msap-policy allowing similar provisioning behavior for automatically created managed SAPs.

17.5 CFHP Policer Control Policy

Provisioning CFHP requires creating policer control policies (policer-control-policy), and applying a policer control policy to the ingress or egress context of a SAP or to the ingress or egress context of a subscriber profile (sub-profile), much the same way scheduler policies (scheduler-policy) are applied.

Applying a policer control policy to a SAP creates an instance of the policy that is used to control the bandwidth associated with the child policers on the SAP. In a similar fashion, an instance of the policy is created when a subscriber profile associated with the policy is applied to a subscriber context. The subscriber policy instance is used to control the bandwidth of the child policers created by the SLA profile instances within the subscriber context.

Policer control policies can only be applied to SAPs created on Ethernet ports. When the policy instance is created, any policers created on the SAP that have an appropriate parent command defined are considered child policers.

17.5.1 Policer Control Policy Root Arbiter

Similar to a scheduler context within a scheduler-policy, the policer-control-policy contains objects called arbiters that control the amount of bandwidth that may be distributed to a set of child policers. Each policer control policy always contains a root arbiter that represents the parent policer. The max-rate defined for the arbiter specifies the decrement rate for the parent policer that governs the overall aggregate rate of every child policer associated with the policy instance. The root arbiter also contains the parent policers MBS configuration parameters that the system uses to individually configure the priority thresholds for each policer instance.

Child policers may parent directly to the root arbiter or to one of the tier 1 or tier 2 explicitly created arbiters.

Each arbiter provides bandwidth to its children using eight strict levels. Children parented at level 8 are first to receive bandwidth. The arbiter continues to distribute bandwidth until either all of its children's bandwidth requirements are met or until the bandwidth its allowed to distribute is exhausted. The root arbiter is special in that its strict priority levels directly represent the priority thresholds within the parent policer.

17.5.2 Tier 1 and Tier 2 Explicit Arbiters

Other arbiters may be explicitly created in the policy for the purpose of creating an arbitrary bandwidth distribution hierarchy. The explicitly created arbiters must be defined within tier 1 or tier 2 on the policy. Tier 1 arbiters must always be parented by the root arbiter and thus become a child of the root arbiter. Any child policers directly parented by a tier 1 policer treat the root arbiter as its grandparent. Inversely, the root arbiter considers the child policers as grandchildren. All grandchild policers inherit the priority level of their parent arbiter (the level that the tier 1 arbiter attaches to the root arbiter) within the parent policer.

An arbiter created on tier 2 may be parented by either an arbiter in tier 1 or by the root arbiter. If the tier 2 arbiter is parented by the root arbiter, it is internally treated the same as a tier 1 arbiter and its child policers have a grandchild to grandparent association with the root arbiter.

When a tier 2 arbiter is parented by a tier 1 arbiter, the child policers parented by a tier 2 arbiter are in a great-grandchild to great-grandparent association with the root arbiter. A great-grandchild policer inherits its indirectly parented tier 1 arbiter's level association with the root arbiter and thus the parent policer.

A child policer's priority level on the root arbiter (directly or indirectly) defines which priority level discards thresholds will be associated with packets mapped to the child policer for use in the parent policer (assuming the packet is not discarded by its child policer).

17.5.3 Explicit Arbiter Rate Limits

The bandwidth a tier 1 or tier 2 arbiter receives from its parent may be limited by the use of the rate command within the arbiter. When a rate limit is defined for a root arbiter, the system enforces the aggregate rate by calculating a per child policer PIR rate based on the distributed bandwidth per child. This calculated PIR is used to override the child's defined PIR and is represented as the child's operational PIR. The calculated rate will never be greater than a child policer's provisioned rate.

17.5.4 CFHP with Child Policer Exceed PIR Enabled

A child policer parented to an arbiter can be enabled to forward traffic exceeding its PIR, in which case:

- Traffic exceeding the operational PIR of the child policer is reprofiled to be exceed-profile, where the operational PIR is determined by the H-pol algorithm from the configuration of the policer parent and the associated arbiters (root and/or intermediate).
- Traffic exceeding the child policer's operational PIR and exceed-profile traffic entering the child policer does not consume capacity from the parent policer (meaning that it does not contribute to the parent policer bucket depth with respect to any of its thresholds).
- Traffic that did not exceed the child policer's operational PIR (when that child is configured with **enable-exceed-pir**) can exceed its parent rate (**max-rate** for the root arbiter) in which case the traffic is forwarded and reprofiled to be exceed-profile and its effect on the child policer is revoked (meaning that it does not contribute to any of the child policer bucket (PIR, CIR, FIR) depths with respect to any of its thresholds).

17.6 CFHP Child Policer Definition and Creation

Policers are created within the context of SAP ingress (sap-ingress) and SAP egress (sap-egress) QoS policies. Policer creation in a QoS policy is defined similar to SAP-based queues. A policer is identified using a policer ID. Queues and policers have different ID spaces (both a policer and queue may be defined with ID 1).

The only create time parameter currently available is the unique policer ID within the policy. Policers do not have a scheduling mode (expedite or best-effort), they also do not need to be placed in in-profile-mode in order to accept traffic from profile in or profile out forwarding classes or subclasses.

All policers within a SAP ingress or egress QoS policy must be explicitly created. No policers are created by default. After a policer is created, forwarding classes or subclasses may be mapped to the policer within the policy. For ingress, each of the individual forwarding types (unicast, multicast, broadcast, and unknown) may be selectively mapped to a policer, policy-created queue or to an ingress port queue group queue. At egress, forwarding classes are not divided into forwarding types, so all packets matched to the forwarding class may be mapped to either a policer, policy-created queue or egress port queue group queue.

Similar to queues, a policer is not created on the SAPs where the policy is applied until at least one forwarding class is mapped to the policer. When the last forwarding class is unmapped from the policer, all the instances of the policer on the SAPs to which the policy is applied are removed.

17.7 Policer Enabled SAP QoS Policy Applicability

Policers are not created on a SAP or subscriber or multiservice site context until at least one forwarding class has been mapped to the policer. Creating a policer within a QoS policy does not cause policers to be created on the SAPs or subscribers or multiservice sites where the policy is applied.

SAP QoS policy applicability and policy policer forwarding class mappings are dependent on policer resource availability. Attempting to map the first forwarding class to a policer causes the policer to be created on the SAPs or subscribers or multiservice site where the policy is applied. If the forwarding plane where the SAP or subscriber or multiservice site exists either does not support policers or has insufficient resources to create the policer for the object, the forwarding class mapping will fail.

When a forwarding class is successfully mapped to a policer within the policy, attempting to apply the policy to a SAP or a subscriber or multiservice site where the policer cannot be created either due to lack of policer support or insufficient policer resources will fail.

Policing is supported only on Ethernet SAPs or Ethernet-based subscribers. Policing is also only supported on FlexPath2-based systems or IOMs with the exception of CCAG and HSMDA SAPs or subscribers. This feature applies to the 7450 ESS and 7750 SR only.

17.8 Child Policer Parent Association

Each policer configured within a SAP ingress or SAP egress QoS policy may be configured to be child policer by defining a parent arbiter association using the parent command. If the command is not executed, the policer operates as a stand-alone policer wherever the policy is applied. If the parent command is executed, but the defined arbiter name does not exist within the SAP context or a subscriber or multiservice site context, the policer is treated as an orphan. The SAP or subscriber or multiservice site context is placed into a degraded state. The system indicates the degraded state by the system setting the ingress-policer-mismatch or egress-policer-mismatch flag for the object. An orphaned policer functions in the same manner as a policer without a parent defined.

An arbiter exists on a SAP when a policer-control-policy containing the arbiter is applied to the appropriate direction (ingress or egress) of the SAP. An arbiter exists on a subscriber when a policer-control-policy containing the arbiter is applied to the subscriber's sub-profile in the appropriate direction as well (applies to the 7450 ESS and 7750 SR).

17.9 Profile-Capped Policers

Profile-capped mode enforces an overall inplus-profile and in-profile burst limit to the CIR bucket for the following packet types:

- ingress undefined
- ingress explicit in-profile
- egress soft-in-profile
- egress explicit inplus
- egress explicit in-profile

The default behavior when profile-capped mode is not enabled is to ignore the CIR output state when an explicit inplus-profile (egress only) and in-profile packet is handled by an ingress or egress policer. The explicit in-profile packets will consume CIR tokens up to two times the CBS at which point the bucket stops incrementing and the CIR output for that type of packet enters the non-conforming state. However, the non-conforming state is ignored by the forwarding plane and the packet continues to be handled as inplus-profile or in-profile. Therefore, the total amount of inplus-profile or in-profile traffic can be greater than the configured CIR.

The profile-capped mode makes two changes.

- At egress, soft-in-profile packets (packets received at ingress as in-profile) are treated the same as explicit in-profile packets (unless explicitly reclassified as out-of-profile) and have an initial policer state of in-profile.
- At both ingress and egress, any packets 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).

A profile-capped policer trusts the in-profile state determined at ingress classification or egress reclassification, so the initial in-profile traffic is preferentially handled with the CIR bucket (two times the CBS instead of CBS used by undefined or soft-out-of-profile traffic) and the total amount of inplus-profile or in-profile traffic output by the policer cannot exceed the CIR (including initial in-profile traffic).

Profile-capped mode has an effect on stat-mode behavior. Each stat mode has a fixed number of counters in the forwarding plane. The mapping of packets to a counter is also fixed by the offered packet state (profile inplus, profile in, profile out, undefined, soft-in-profile and soft-out-of-profile) in conjunction with the output state of the policer. In the non-capped mode, soft-in-profile is considered undefined while in capped mode it is considered to be equivalent to profile in. Another potential issue with ingress and egress stat-modes is the fact that green packets (that is, those that are profile in at ingress and egress, or soft-in-profile at egress) can turn yellow in the policer output.

[Table 129](#) describes how the CIR rate and initial profile of each packet affects the output of normal (non-profile-capped) and profile-capped mode policers.

Table 129 Effect of Profile-Capped Mode on CIR Output

CIR Setting	Initial Profile State	Normal Mode	Profile-Capped Mode	Notes
CIR=0	Ingress undefined	Always out-of-profile	Always out-of-profile	When CIR = 0, the CIR has no effect on the packet profile except for ingress-undefined packets. If profile-capped is used, this forces all packets to be out-of-profile except for those explicitly reprofiled to exceed-profile.
	Ingress profile in	Always in-profile	Always out-of-profile	
	Ingress profile out	Always out-of-profile	Always out-of-profile	
	Egress soft-in-profile	Always in-profile	Always out-of-profile	
	Egress soft-out-of-profile	Always out-of-profile	Always out-of-profile	
	Egress profile inplus	Always inplus-profile	Always out-of-profile	
	Egress profile in	Always in-profile	Always out-of-profile	
	Egress profile out	Always out-of-profile	Always out-of-profile	
	Egress profile exceed	Always exceed-profile	Always exceed-profile	

Table 129 Effect of Profile-Capped Mode on CIR Output (Continued)

CIR Setting	Initial Profile State	Normal Mode	Profile-Capped Mode	Notes
CIR=Max/ PIR	Ingress undefined	Always in-profile	Always in-profile	CIR never reaches non-conforming state.
	Ingress profile in	Always in-profile	Always in-profile	
	Ingress profile out	Always out-of-profile	Always out-of-profile	
	Egress soft-in-profile	Always in-profile	Always in-profile	
	Egress soft-out-of-profile	Always in-profile	Always in-profile	
	Egress profile inplus	Always inplus-profile	Always inplus-profile	
	Egress profile in	Always in-profile	Always in-profile	
	Egress profile out	Always out-of-profile	Always out-of-profile	
	Egress profile exceed	Always exceed-profile	Always exceed-profile	

Table 129 Effect of Profile-Capped Mode on CIR Output (Continued)

CIR Setting	Initial Profile State	Normal Mode	Profile-Capped Mode	Notes
0 < CIR < PIR	Ingress undefined	In-profile below CBS Out-of-profile at or above CBS	In-profile below CBS Out-of-profile at or above CBS	
	Ingress profile in	Always in-profile	In-profile below two times CBS Out-of-profile at or above two times CBS	
	Ingress profile out	Always out-of-profile	Always out-of-profile	
	Egress soft-in-profile	In-profile below CBS Out-of-profile at or above CBS	In-profile below two times CBS Out-of-profile at or above two times CBS	
	Egress soft-out-of-profile	In-profile below CBS Out-of-profile at or above CBS	In-profile below CBS Out-of-profile at or above CBS	
	Egress profile inplus	Always inplus-profile	Inplus-profile below two times CBS Out-of-profile at or above two times CBS	
	Egress profile in	Always in-profile	In-profile below two times CBS Out-of-profile at or above two times CBS	
	Egress Profile Out	Always Out-of-profile	Always Out-of-profile	
	Egress Profile Exceed	Always Exceed-profile	Always Exceed-profile	

17.10 Policer Interaction with Profile, Discard Eligibility, and Ingress Priority

Packets that are offered to an ingress policer may have three different states relative to initial profile:

- **undefined** — Either the forwarding class or subclass associated with the packet is not explicitly configured as profile in; profile out or de-1-out-profile is enabled and the dot1p DE bit is set to zero.
- **in-profile** — The forwarding class or subclass associated with the packet is configured as profile in.
- **out-of-profile** — The forwarding class or subclass associated with the packet is configured as profile out or de-1-out-profile is enabled, and the dot1p DE bit is set to 1.

Ingress policed packets are not subject to ingress queue CIR profiling within the ingress policer output queues. While the unicast and multipoint shared queues used by the system for ingress queuing of policed packets may have a CIR rate defined, this CIR rate is only used for rate-based dynamic priority scheduling purposes. The state of the CIR bucket while forwarding a packet from a policer-output-queues shared queue will not alter the packets ingress in-profile or out-of-profile state derived from the ingress policer.

Priority high and low are used in the child policer's PIR leaky bucket to choose one of two discard thresholds (threshold-be-low and threshold-be-high) that are derived from the child policer's mbs and high-priority-only parameters. The high threshold is directly generated by the mbs value. The low threshold is generated by reducing the mbs value by the high-priority-only percentage. A packet's priority is determined while the packet is evaluated against the ingress classification rules in the sap-ingress QoS policy.

Packets that are offered to an egress policer may have six different states relative to their initial profile:

- **soft-in-profile**—the final result at ingress was in-profile and the packet's profile has not been reclassified at egress
- **soft-out-of-profile**—the final result at ingress was out-of-profile and the packet's profile has not been reclassified at egress
- **hard-inplus-profile**—the profile of the packet has been reclassified at egress as profile inplus
- **hard-in-profile**—the profile of the packet has been reclassified at egress as profile in

- hard-out-of-profile—the profile of the packet has been reclassified at egress as profile out
- hard-exceed-profile—the profile of the packet has been reclassified at egress as profile exceed

When an egress policer's CIR rate is set to 0 (or not defined), the policer has no effect on the profile of packets offered to the policer. An exception to this is when **enable-exceed-pir** is configured under the policer. In this case, the exceed-profile state of a packet takes precedence over the hard-out/in/inplus reclassification, specifically, traffic that is reprofiled to exceed within a SAP egress policer has an exceed-profile state regardless of whether it was reclassified at egress to hard-out, hard-in, or hard-inplus.

Setting a non-zero rate for the egress policer's CIR will modify this behavior for DSCP, IP precedence, dot1p, and DEI egress marking purposes. Hard-inplus-profile and hard-in-profile retain their inherent inplus-profile or in-profile behavior and the hard-out-of-profile and hard-exceed-profile retain their inherent out-of-profile or exceed-profile behavior.

When the egress packet state is soft-in-profile and soft-out-of-profile and the policer's CIR is configured as non-zero, the current CIR state of the policer's CIR bucket will override the packet's soft profile state. When the policer's CIR is currently conforming, the output will be in-profile. When the CIR state is currently exceeding, the output will be out-of-profile.

Hard-exceed-profile packets are discarded by default by an egress policer. If **enable-exceed-pir** is configured, the hard-exceed-profile packets are forwarded and, when the PIR state is exceeding, all packets are forwarded with an exceed-profile state.

For egress marking decisions, the hard-inplus-profile, hard-in-profile, and hard-out-of-profile packets ignore the egress policer's CIR state. When the packet state is hard-inplus-profile or hard-in-profile, the in-profile dot1p marking will be used, and when DEI marking is enabled for the packet's forwarding class, the exceed-profile traffic will be marked 0. When the packet state is hard-out-of-profile or hard-exceed-profile, the out-of-profile dot1p marking will be used, unless explicit dot1p exceed-profile marking is configured, in which case the exceed-profile traffic will be marked with the configured value, and when DEI marking is enabled for the packets forwarding class, the exceed-profile traffic will be marked 1.

The dot1p, outerDot1p, and DEI (when DE marking is configured) will reflect the CIR- and PIR-derived packet state. If the **enable-dscp-prec-remark** command is enabled, the DSCP and IP precedence will reflect the CIR- and PIR-derived packet state.

17.10.1 Ingress ‘Undefined’ Initial Profile

Access ingress packets have one of three initial profile states prior to processing by the policer:

- undefined
- profile in
- profile out

The SAP ingress QoS policy classification rules map each packet to either a forwarding class or a subclass within a forwarding class. The forwarding class or subclass may be defined as explicit profile in or profile out (the default is no profile). When a packet’s forwarding class or subclass is explicitly defined as profile in or profile out, the packet’s priority is ignored, and it is not handled by the ingress policer as profile ‘undefined’.

See [Table 129](#) to track the ingress behavior of initial profile and the effect of the CIR bucket on that initial state.

At egress, an ingress policer output of ‘in-profile’ is treated as ‘soft-in-profile’ and an ingress policer output of ‘out-of-profile’ is treated as ‘soft-out-of-profile’. Each may be changed by egress profile reclassification or by an egress policer with a CIR rate defined.

17.10.2 Ingress Explicitly ‘In-Profile’ State Packet Handling without Profile-Capped Mode

Packets that are explicitly ‘in-profile’ remain ‘in-profile’ in the ingress forwarding plane and are not affected by the ingress policer CIR bucket state when profile-capped mode is not enabled. They do not bypass the policer’s CIR leaky bucket but are extended with a greater threshold than the CBS derived ‘threshold-bc’. This allows the ‘undefined’ packets to backfill the remaining conforming CIR bandwidth after accounting for the explicit ‘in-profile’ packets. This does not prevent the sum of the explicit ‘in-profile’ from exceeding the configured CIR rate, but it does cause the ‘undefined’ packets that are marked ‘in-profile’ to diminish to zero when the combined explicit ‘in-profile’ rate and ‘undefined’ rate causes the bucket to reach ‘threshold-bc’.

The policer's CIR bucket will indicate that the explicit 'in-profile' packets should be marked 'out-of-profile' when the bucket reaches the greater threshold, but this indication is ignored by the ingress forwarding plane. All explicit 'in-profile' packets remain in-profile within the ingress forwarding plane. However, when the packet is received at egress, an ingress 'in-profile' packet will be treated as 'soft-in-profile' and the profile may be changed either by explicit profile reclassification or by an egress policer with a CIR rate defined.

Explicit in-profile packets do not automatically use the high-priority threshold ('threshold-be-high') within the child policer's PIR bucket. If preferential burst tolerance is desired for explicit in-profile packets, the packets should also be classified as priority high.

17.10.3 Ingress Explicitly 'In-Profile' State Packet Handling with Profile-Capped Mode

When profile-capped mode is enabled, the packet handling behavior defined in [Ingress 'Undefined' Initial Profile](#) is altered in one aspect. The CIR output state of yellow at the greater threshold is actually honored and the packet will be treated as out-of-profile. The packet will be sent to egress in the 'soft-out-of-profile' state in this case.

17.10.4 Ingress Explicit 'Out-of-Profile' State Packet Handling

Packets that are explicitly 'out-of-profile' remain 'out-of-profile' in the ingress forwarding plane. Unlike initially 'in-profile' packets, they do not consume the policer's CIR bucket depth (accomplished by setting the 'threshold-bc' to 0) and thus do not have an impact on the amount of 'undefined' marked as 'in-profile' by the policer.

While explicit 'out-of-profile' packets remain out-of-profile within the ingress forwarding plane, the egress forwarding plane treats ingress out-of-profile packets as 'soft-out-of-profile' and the profile may be changed either by explicit profile reclassification or by an egress policer with a CIR rate defined.

17.10.5 Egress Explicit Profile Reclassification

An egress profile reclassification overrides the ingress-derived profile of a packet and may set it to hard-inplus-profile, hard-in-profile, hard-out-of-profile, or hard-exceed-profile. A packet that has not been reclassified at egress retains its soft-in-profile or soft-out-of-profile status.

Egress inplus-profile and in-profile (including soft-in-profile and hard-in-profile) packets use the child policer's high threshold-be value within the child policer's PIR bucket while soft-out-of-profile and hard-out-of-profile packets use the child policer's low threshold-be value. Egress hard-exceed-profile packets are not subject to any threshold control in the child's PIR bucket if **enable-exceed-pir** is configured; otherwise, they are discarded.

17.10.6 Preserving Out of Profile State at Egress Policer

Traffic sent through an egress policer with a non-zero CIR will be reprofiled by default based on the CIR threshold of the egress policer. To accommodate designs where traffic is set to be out-of-profile at ingress, and the out-of-profile state is required to be maintained by an egress policer, the parameter **profile-out-preserve** can be configured under the egress policer. Explicit egress reclassification to the profile takes precedence over the profile-out-preserve operation.

17.10.7 Egress Policer CIR Packet Handling without Profile-capped Mode

When an egress policer has been configured with a CIR (maximum or explicit rate other than 0) and profile-capped mode is not enabled, the policer's CIR bucket state overrides the ingress soft-in-profile or soft-out-of-profile state much like the ingress policer handles initial profile undefined packets. If the CIR has not been defined or set to 0 on the egress policer, the egress policer output state will be in-profile for soft-in-profile packets and out-of-profile for soft-out-of-profile packets.

If a packet's profile has been reclassified at egress, the new profile classification is handled in the same way as the ingress policer handling of initial in-profile or out-of-profile packets. When a packet has been reclassified as hard-inplus-profile or hard-in-profile, it is applied to the egress policer's CIR bucket using a threshold-bc higher than the threshold-bc derived from the policer's CBS parameter, but the policer output profile state will remain inplus-profile or in-profile even if the higher threshold is crossed. When a packet has been reclassified as hard-out-of-profile or hard-exceed-profile, it does not consume the egress policer's CIR bucket depth and the policer output profile state remains out-of-profile or exceed-profile.

17.10.8 Egress Policer CIR Packet Handling with Profile-capped Mode

When profile-capped mode is enabled, the egress packet handling described in [Egress Policer CIR Packet Handling without Profile-capped Mode](#) is modified in three ways.

First, the soft-in-profile packets received from ingress are handled in the same way as egress explicit profile in reclassification, unless the packet has been reclassified to profile out or profile exceed at egress.

Second, explicit egress inplus-profile, in-profile, and soft-in-profile packets that have not been reclassified to out-of-profile or exceed-profile at egress are allowed to be marked out-of-profile by an egress policer with a CIR not set to 0.

Third, when a policer has a CIR set to 0 (the default rate), all profile-capped packets are treated as out-of-profile independent of the initial profile state, except for exceed-profile packets that remain as exceed-profile.

17.10.9 Forwarding Traffic Exceeding PIR in Egress Policers

An egress policer can be configured to forward traffic that enters the policer with an exceed-profile state or exceeds its operational PIR instead of dropping it. The traffic exceeding the PIR is assigned an exceed-profile state. This is supported for any configured (not dynamic) policer in a SAP egress QoS policy, which can be used for both SAPs and subscribers, and in an egress queue group template that can be used on egress network ports.

When **enable-exceed-pir** is configured under the policer, the exceed-profile state of a packet takes precedence over the hard-out/in/inplus reclassification, specifically, traffic that is reprofiled to exceed within a SAP egress policer has an exceed-profile state regardless of whether it was reclassified at egress to hard-out, hard-in, or hard-inplus.

The **stat-mode offered-total-cir-exceed** command provides forward and drop counters for exceed-profile traffic, as follows:

```
configure
  qos
    queue-group-templates
      egress
        queue-group <queue-group-name> create
        policer <policer-id> create
          enable-exceed-pir
          stat-mode offered-total-cir-exceed
        exit
      exit
    exit
  sap-egress <policy-id> create
    policer <policer-id> create
      enable-exceed-pir
      stat-mode offered-total-cir-exceed
    exit
  exit
exit
exit
```

The dot1p, outer dot1p, DSCP, and precedence can be remarked for the exceed-profile traffic.

17.10.10 Post Egress Policer Packet Forwarding Class and Profile State Remapping

Packets processed by a SAP or subscriber egress child policer can have their forwarding class and profile state remapped to a different forwarding class and profile state after exiting the policer. Remapping is achieved by using a post-policer mapping policy containing mapping statements that determine the remapping of packets based on their forwarding class and profile state.

The post-policer mapping policy is configured as follows:

```
configure
  qos
    post-policer-mapping <mapping-policy-name> [create]
      description <description-string>
      fc <fc-name> profile <profile> [create]
```

```
maps-to fc <fc-name> profile <profile>
```

where:

- <mapping-policy-name>: name of mapping policy, up to 32 characters
- <description> : description text string
- <fc-name> : af, be, ef, h1, h2, l1, l2, nc
- <profile> : in, out, exceed, inplus

Up to 32 mapping statements are supported within a policy, covering the maximum combinations of eight forwarding classes and four profile states. A maximum of seven post-policer mapping policies can be configured per system.

After being configured, the post-policer mapping policy must be applied within a SAP egress QoS policy with, at most, one mapping policy per SAP egress QoS policy:

```
configure
  qos
    sap-egress <policy-id> [create]
      post-policer-mapping <mapping-policy-name>
```

Packets entering a child policer are assigned a forwarding class and profile state. The configuration of the policer can change the profile state of the exiting packet, but not the forwarding class. If a post-policer mapping policy is applied within the SAP egress QoS policy, a packet exiting the policer with a given forwarding class and profile state can be remapped to a different forwarding class and profile state. For example, consider the following post-policer mapping policy:

```
configure
  qos
    post-policer-mapping ppm1 create
      fc ef profile exceed create
      maps-to fc be profile out
    exit
  exit
```

Packets exiting an egress child policer with forwarding class “ef” and profile state “exceed” will have their forwarding class remapped to “be” and their profile state to “out”.

The mapping applies to all policers within the SAP egress QoS policy, including regular child policers and policers configured in an IP/IPv6 criterion action statement, except for dynamic policers. The remapping does not affect the policer statistics or the parent policer processing (root arbiter) as it occurs after each of these.

The new forwarding class is used to select the egress queue on which the post-policer traffic is placed. The new profile is used to determine the congestion control handling in that queue and its pool, specifically the drop tail or slope that is applied.

Egress packet remarking is based on the marking configured for the forwarding class and profile of the traffic after being policed but before it is remapped.

This remapping is supported on platforms using FP3- and higher-based hardware, with the exception of the 7750 SR-a4/a8, which does not support egress policers.

Remapping a subset of packets from one forwarding class to another could result in out-of-order packets being received at the destination if there is congestion or different latency characteristics on the paths of the different forwarding classes.

17.10.11 Ingress Child Policer Stat-Mode

A policer has multiple types of input traffic and multiple possible output states for each input traffic type. These variations differ between ingress and egress.

For ingress policing, each offered packet has a priority and a profile state. The priority is used by the policer to choose either the high- or low-priority PIR threshold-be. Every offered packet is either priority high or priority low. The offered profile state defines how a packet will interact with the policers CIR bucket state. The combinations of priority and initial profile are as follows:

- Offered priority low, undefined profile
- Offered priority low, explicit profile in
- Offered priority low, explicit profile out
- Offered priority high, undefined profile
- Offered priority high, explicit profile in
- Offered priority high, explicit profile out



Note: When de-1-out-profile is enabled, DEI = 0 is considered as undefined profile and DEI = 1 is considered the same as profile out.

The possible output results for the ingress policer are:

- Output green (in-profile)
- Output yellow (out-of-profile)
- Output red (discard)

In order to conserve counter resources, the system supports a policer **stat-mode** command that is used to identify what counters are actually needed for the policer. Not every policer will have a CIR defined, so the output green/yellow states will not exist. Also, not every policer will have both high- and low-priority or explicit in-profile or out-of-profile offered traffic types. Essentially, the **stat-mode** command allows the counter resources to be allocated based on the accounting needs of the individual policers.

Setting the **stat-mode** does not modify the packet handling behavior of the policer. For example, if the configured stat-mode does not support in-profile and out-of-profile output accounting, the policer is not blocked from having a configured CIR rate. The CIR rate will be enforced, but the amount of in-profile and out-of-profile traffic output from the policer will not be counted separately (or maybe not at all based on the configured stat-mode).

A policer is created with minimal counters sufficient to provide total offered and total discarded (the total forwarded is computed as the sum of the offered and discarded counters). The **stat-mode** is defined within the **sap-ingress** or **sap-egress** QoS policy in the policer context. When defining the **stat-mode**, the counter resources needed to implement the mode must be available on all forwarding planes where the policer has been created using the QoS policy unless the policer instance has a stat-mode override defined. Use the **tools dump resource-usage card fp** command to see the resources used and available. If insufficient resources exist, the change in the mode will fail without any change to the existing counters currently applied to the existing policers. If the QoS policy is being applied to a SAP or subscriber or multiservice site context and insufficient counter resources exist to implement the configured modes for the policers within the policy, the QoS policy will not be applied. For SAPs, this means the previous QoS policy will stay in effect. For subscribers, it could mean that the subscriber host event where the QoS policy is being applied will fail and the subscriber host may be blocked or removed.

A stat-mode with at least minimal stats is required before the policer can be assigned to a parent arbiter using the parent command.

Successfully changing the stat-mode for a policer causes the counters associated with the policer to reset to zero. Any collected stats on the object the policer is created on will also reset to zero.

The system uses the forwarding plane counters to generate accounting statistics and for calculating the operational PIR and FIR rates for a set of children when they are managed by a policer-control-policy. Only the offered counters are used in hierarchical policing rate management. When multiple offered stats are maintained for a child policer, they are summed to derive the total offered rate for each child policer.

All ingress policers have a default CIR value of 0, meaning that by default, all packets except packets classified as profile in will be output by the policer as out-of-profile. This may have a negative impact on egress marking decisions (if in-profile and out-of-profile have different marking values) and on queue congestion handling (WRED or queue drop tail decisions when out-of-profile is less preferred). The following options exist to address this potential issue:

- If all packets handled by the policer must be output as in-profile by the policer, either the packet's forwarding class or subclass can be defined as profile in or the CIR on the policer can be defined as max
- If some packets must be output as in-profile while others output as out-of-profile, three options exist:
 - The CIR may be left at '0' while mapping the packets that must be output as in-profile to a forwarding class or subclass provisioned as profile in
 - The CIR may be set to max while mapping the packets that must be output as out-of-profile to a forwarding class or subclass provisioned as profile out
 - Ignore the CIR on the policer and solely rely on the forwarding class or subclass profile provisioning to the proper policer CIR output

Make sure to use the correct stat mode if the policer's CIR is explicitly not set or is set to 0. The **no-cir** version of the stat-mode must be used when the CIR has a non-zero value. Also, when overriding the policer's CIR mode, make sure to override the stat-mode instance (CIR override can be performed using SNMP access).

Ingress supported stat-modes are:

- no-stats
- minimal - default
- 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

17.10.12 Egress Child Policer Stat-Mode

All packets received on the egress forwarding plane are profiled as either in-profile or out-of-profile. The egress forwarding plane treats the ingress-derived profile as a soft state that may be either overridden by an egress profile reclassification or by a CIR rate enforced by an egress policer.

Egress policers have a default CIR set to 0, but in the egress case a value of 0 disables policer profiling. Egress packets on a CIR-disabled egress policer retain their offered profile state (soft-in-profile, soft-out-of-profile, hard-inplus, hard-in-profile, hard-out-of-profile, or hard-exceed-profile) unless the **enable-exceed-pir** command is configured, in which case the exceed-profile state of a packet takes precedence over the hard-out/in/inplus reclassification.

For egress, the possible types of offered packets include:

- Soft offered in-profile (from ingress)
- Soft offered out-of-profile (from ingress)
- Egress explicit inplus-profile (reclassified at egress)
- Egress explicit in-profile (reclassified at egress)
- Egress explicit out-of-profile (reclassified at egress)
- Egress explicit exceed-profile (reclassified at egress)

The possible output results are:

- Output inplus-profile
- Output in-profile
- Output out-of-profile
- Output discard or exceed-profile

The stat-mode command follows the same counter resource rules as ingress.

Egress supported stat-modes are:

- no-stats
- minimal - default
- offered-profile-no-cir
- offered-profile-cir
- offered-total-cir
- offered-limited-capped-cir
- offered-profile-capped-cir

- offered-total-cir-exceed
- offered-four-profile-no-cir
- offered-total-cir-four-profile

Details of the output showing the stat-modes for ingress and egress child policers are in the Class Fair Hierarchical Policing for SAPs section of the *SR OS Advanced Configuration Guide*.

17.11 Profile-Preferred Mode Root Policers

Configuring the **profile-preferred** option gives preference for inplus-profile packets or in-profile packets to consume the root policer PIR bucket tokens at a given priority level. This preference applies to packets whose profile is either configured explicitly or set by the output of the child policer CIR bucket.

When this option is selected, all child policers parented to a root policer will have their FIR bucket track the state of the CIR bucket. In other words, an inplus-profile or in-profile packet will always be blue and an out-of-profile packet will always be orange. When admitting packets from the child policers within a given priority level, orange packets will be allowed up to the discard-unfair threshold while blue packets will be allowed up to the discard-all threshold. If a child policer is configured to forward traffic exceeding its PIR, the exceed-profile traffic does not contribute to the parent policer bucket depth with respect to any of its thresholds.

The **profile-preferred** option forces the FIR bucket to track the CIR bucket's decrement rate and the threshold chosen for the CIR bucket is also be used in the FIR bucket (instead of using the threshold associated with the PIR bucket).

The inplus/in/out/exceed-profile output from the policer is used for packet marking decisions. The blue/orange child policer input to the parent policer chooses the discard-orange or discard-all thresholds for the child policer's priority level within the parent policer.

Explicit inplus-profile and in-profile packets stay blue up to the high CBS threshold, undefined profile packets stay blue up to the low CBS threshold (1x CBS) and explicit out-of-profile packets are always orange due to a 0 CBS threshold. Orange packets are discarded by the parent policer within the child policer's priority level before the blue packets, preferring blue packets over orange when the discard-orange threshold is crossed.

Use the following CLI syntax to configure the **profile-preferred** option. This option also applies to overrides applied to instances of a policer control policy under a SAP or subscriber or multiservice site context.

```
config qos
  policer-control-policy policy-name [create]
  no policer-control-policy policy-name
  description "description-string"
  no description
  root
    max-rate {kilobits-per-second | max}
    no max-rate
    [no] profile-preferred
    priority-mbs-thresholds
      min-thresh-separation size [bytes | kilobytes]
      no min-thresh-separation
```

```
priority level
  mbs-contribution size [bytes | kilobytes] [fixed]
  no mbs-contribution
```

The profile-preferred option provides us a way to configure a specific FIR (since it uses the CIR as FIR). In the direct-parented case (no intermediate arbiters present at all) the child policers do not need to have their offered rate polled as each policer will always have PIR equal to the min (child PIR, root PIR) and the FIR and CIR are fixed and equal. The child parenting weights are thus not used. This impacts the show commands, for example, offered rate information will not be available. The output of some show commands (**show qos policer-hierarchy ... detail**) should be adjusted for profile-preferred configurations.

If an intermediate arbiter is present, then polling is offered at different rates since the child policer PIRs will be set based on this information so as to share the intermediate arbiter PIR in proportional to their parenting weight to the intermediate arbiter.

17.12 Child Policer Hierarchical QoS Parenting

Policers can be parented into the QoS hierarchy used for queue and scheduler bandwidth control, referred to as hierarchical QoS (HQoS). This allows the bandwidth of policers, queues, and schedulers to be managed in the same HQoS hierarchy.

HQoS builds a scheduling hierarchy for a queue by parenting it into a scheduler or port scheduler. The schedulers can be parented into other schedulers to create multiple tiers or into a port scheduler that can exist on a Vport or port.

Parenting a policer into HQoS is supported at egress for subscribers and SAPs, with multiservice sites (MSS) supported for SAPs. A post-policer local queue is not supported with HQoS managed policers, nor are queues that are mapped by the **use-fc-mapped-queue** parameter in a criteria action statement. Parenting a policer into HQoS is not supported on the 7950 XRS.

To enable the parenting of an egress policer into HQoS, the following command is configured in the SAP egress QoS policy:

CLI Syntax:

```
configure
  qos
    sap-egress policy-id
      policers-hqos-manageable
    exit
```

When the **policers-hqos-manageable** command is configured, all policers in the SAP egress QoS policy, except for dynamic policers, can be managed by HQoS. To be managed by HQoS, the policer must be configured with either a **scheduler-parent** or **port-parent** command, or be orphaned to an egress port scheduler applied on a Vport or port.

The **no policers-hqos-manageable** command results in policers not being managed by HQoS.

If the **policers-hqos-manageable** command is used, the **parent-location sla**, **policers enable-exceed-pir**, or **policers stat-mode no-stats** commands may not be used within an SAP egress QoS policy.

Egress policers can be parented into a scheduler in a scheduler policy using the **scheduler-parent** command:

CLI Syntax:

```
configure
  qos
    sap-egress policy-id
      policer policer-id
```



```

scheduler-parent scheduler-name
                    [weight weight] [level level]
                    [cir-weight cir-weight] [cir-
                    level cir-level]
exit
exit

```

When a scheduler is specified, no checks are performed as to whether the scheduler exists. If the *scheduler-name* does not exist, the policer is placed in an orphaned operational state. The policer will accept packets but will not be bandwidth-limited by a virtual scheduler or the scheduler hierarchy applied to the SAP or subscriber. Consequently, an orphan policer operates in the same way as a non-HQoS-managed policer. On a SAP, the orphan state is indicated in the **show service sap-using** command output with the SapEgressPolicerMismatch flag. This flag is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The **level** and **cir-level** keywords define the level in the HQoS hierarchy to which the policer will parent for the above-CIR and within-CIR bandwidth distribution passes, respectively. If the **cir-level** is set to 0, the policer will not get any bandwidth allocated in the within-CIR pass.

The **weight** and **cir-weight** keywords define the relative weight of this policer in comparison with other child policers, queues, or schedulers at the same level when competing for bandwidth on the parent *scheduler-name* at the above-CIR and within-CIR bandwidth distribution passes, respectively. If the **weight** or **cir-weight** is set to 0, the policer receives bandwidth only after other children with a non-zero weight at this level are serviced.

If **limit-unused-bandwidth** is configured in the HQoS hierarchy to which the policer is parented, only the offered rate increasing in the last sampled period is used to determine that the policer has accumulated work.

Egress policers can also be parented into a port scheduler using the **port-parent** command:

CLI Syntax:

```

configure
  qos
    sap-egress policy-id
      policer policer-id
        port-parent [weight weight] [level
                    level] [cir-weight cir-weight]
                    [cir-level cir-level]

```

If the exit port used by the policed traffic is configured with a port scheduler but the policer has neither a scheduler parent nor a port parent, or if it is orphaned (its scheduler parent does not exist), then the policer will be fostered by the port scheduler.

When parenting to a port scheduler, the subscriber profile or SAP can use an egress aggregate rate limit to control its traffic rate.

The policer **scheduler-parent** and **port-parent** commands are mutually exclusive; configuring one will override the other. These commands and with the policer **parent** command (that parents the policer to an arbiter) are also mutually exclusive.

The system does not prevent the configuration of a policer control policy for a SAP, multiservice site, or subscriber using HQoS managed policers. The arbiters for these policer control policies are not used but are allocated, so must be accounted for when considering scaling.

The configuration of **profile-out-preserve** and **profile-capped** is permitted for HQoS policers with these configurations affecting the within-CIR and above-CIR statistics for the HQoS managed policer.

The purpose of parenting a policer to HQoS is to measure the policer traffic in the HQoS hierarchy so that the configured HQoS bandwidth allocation can be enforced. When traffic passes through the policer, it exits through an access egress queue group queue. If the queue group queue was also parented to the same HQoS hierarchy, the policed traffic would be measured twice: once through the HQoS managed policer, then a second time through a post-policer access egress queue group queue. To prevent the traffic from being measured the second time, the queue group queues must be configured so that they are not managed by HQoS, as follows:

CLI Syntax:

```
configure
  qos
    queue-group-templates
      egress
        queue-group queue-group-name
        no queues-hqos-manageable
```

The default for the **queues-hqos-manageable** command is to allow the queues to be managed by HQoS.

When **no queues-hqos-manageable** is configured, all queues in the egress queue group instances using this template are not managed by HQoS. This command and the configuration of policers and queue **packet-byte-offset** within the egress queue group template are mutually exclusive. The configuration of **no queues-hqos-manageable** is permitted in the default egress **policer-output-queue** queue group template, which avoids the need to create additional queue groups when policers managed by HQoS are used.

When a queue group template with **no queues-hqos-manageable** is configured under a port's Ethernet access egress context, the configuration of an aggregate rate or scheduler policy is not permitted under that context, nor are parent overrides for any of the queues in the queue group. If a port scheduler is configured on the port, the queue group queues are not parented to the port scheduler.

The configuration of an **encap-offset** within the egress of a subscriber profile does not apply to policer traffic that exits through egress queue group queues.

A queue group configured with **no queues-hqos-manageable** should only be used for post-policer traffic from policers in a SAP egress QoS policy configured with **policers-hqos-manageable**. In this case, the traffic is only measured once by HQoS. The following scenarios should be avoided because traffic will either not be measured, or will be measured twice by HQoS, which will cause the HQoS result to be inaccurate:

- traffic through policers in a SAP egress QoS policy configured with **policers-hqos-manageable**, exiting through a queue group queue with its queue group template configured with **queues-hqos-manageable**. In this case, the traffic will be measured twice by HQoS.
- traffic through policers in a SAP egress QoS policy configured with **no policers-hqos-manageable**, exiting through a queue group queue with its queue group template configured with **no queues-hqos-manageable**. In this case, the traffic will not be measured by HQoS at all.
- traffic that is redirected in a SAP egress QoS policy to the queue group queue that has **no queues-hqos-manageable** configured in its queue group template. In this case, the traffic will not be measured by HQoS.

For each of the preceding cases, a mismatch flag, *SapEgressHQoSMgmtMismatch*, is displayed in the **show service sap-using** output. For the first two cases, the **show subscriber-mgmt** output indicates *Egr hqos mgmt status : mismatch* under the SLA profile instance (when the traffic is measured only once, the output displays *Egr hqos mgmt status : enabled*).

17.13 Class Fair Hierarchical Policing (CFHP) Policy Command Reference

17.13.1 Command Hierarchies

17.13.1.1 Class Fair Hierarchical Policing Commands

```

config
  — qos
    — policer-control-policy policy-name [create]
    — no policer-control-policy
      — description description string
      — no description
      — root
        — max-rate {rate | max}
        — no max-rate
        — [no] profile-preferred
        — priority-mbs-thresholds
          — min-thresh-separation [size] [bytes | kilobytes]
          — no min-thresh-separation
          — priority level
            — [no] mbs-contribution
        — tier {1 | 2}
          — arbiter arbiter-name [create]
          — no arbiter arbiter-name
            — description description-string
            — no description
            — parent {root | arbiter-name} [level priority-level] [weight weight-level]
            — no parent
            — rate {kilobits-per-second | max}
            — no rate
          — post-policer-mapping mapping-policy-name [create]
          — no post-policer-mapping mapping-policy-name
            — description description-string
            — no description
            — fc fc-name profile profile [create]
            — no fc fc-name profile profile
              — maps-to fc fc-name profile profile

```

17.13.1.2 Operational Commands

```

config

```

```
— qos
  — copy
    — post-policer-mapping src-name dst-name [overwrite]
```

17.13.1.3 Show Commands

```
show
  — qos
    — post-policer-mapping [post-policer-name] [association | detail]
```

17.13.2 Command Descriptions

17.13.2.1 Configuration Commands

17.13.2.1.1 Generic Commands

description

Syntax	description <i>description string</i> no description
Context	config>qos>policer-control-policy config>qos>post-policer-mapping
Description	The description command is used to define an informational ASCII string associated with the policer control policy or post policer mapping policy. The string value can be defined or changed at any time when the policy exists. The no form of this command is used to remove an explicit description string from the policer.
Default	no description
Parameters	<i>description string</i> — The description-string parameter defines the ASCII description string for the policer control policy. Allowed values are any string of up to 80 characters, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc), the entire string must be enclosed within double quotes. Default None

17.13.2.1.2 Class Fair Hierarchical Policing Commands

policer-control-policy

Syntax	policer-control-policy <i>policy-name</i> [create] no policer-control-policy
Context	config>qos
Description	This command is used to create, delete, or modify policer control policies. The policer-control-policy controls the aggregate bandwidth available to a set of child policers. When created, the policy can be applied to ingress or egress SAPs. The policy can also be applied to the ingress or egress context of a sub-profile.
Default	no policer-control-policy
Parameters	<i>policy-name</i> — Each policer-control-policy must be created with a unique policy name. The <i>policy-name</i> must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode. Default None create — The create keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

root

Syntax	root
Context	config>qos>policer-control-policy
Description	The root node contains the policer control policies configuration parameters for the root arbiter. Within the node, the parent policer's maximum rate limit can be set, the strict priority level, and fair threshold portions may be defined per priority level. The root node always exists and does not need to be created.
Default	None.

max-rate

Syntax	max-rate { <i>kilobits-per-second</i> max }
	no max-rate

Context	config>qos>plcr-ctrl-plcy>root
Description	The max-rate command defines the parent policer's PIR leaky bucket's decrement rate. A parent policer is created for each time the policer-control-policy is applied to either a SAP or subscriber instance for the 7450 ESS and 7750 SR, or multiservice site instance for the 7950 XRS. Packets that are not discarded by the child policers associated with the SAP or subscriber or multiservice site instances are evaluated against the parent policer's PIR leaky bucket. For each packet, the bucket is first decremented by the correct amount based on the decrement rate to derive the current bucket depth. The current depth is then compared to one of two discard thresholds associated with the packet. The first discard threshold (discard-unfair) is applied if the FIR (Fair Information Rate) leaky bucket in the packet's child policer is in the conforming state. The second discard threshold (discard-all) is applied if the child policer's FIR leaky bucket is in the exceed state. Only one of the two thresholds is applied per packet. If the current depth of the parent policer PIR bucket is less than the threshold value, the parent PIR bucket is in the conform state for that particular packet. If the depth is equal to or greater than the applied threshold, the bucket is in the violate state for the packet. If the result is "conform," the bucket depth is increased by the size of the packet (plus or minus the per-packet-offset setting in the child policer) and the packet is not discarded by the parent policer. If the result is "violate," the bucket depth is not increased and the packet is discarded by the parent policer. When the parent policer discards a packet, any bucket depth increases (PIR, CIR and FIR) in the parent policer caused by the packet are canceled. This prevents packets that are discarded by the parent policer from consuming the child policers PIR, CIR, and FIR bandwidth. The policer-control-policy root max-rate setting may be overridden on each SAP or sub-profile where the policy is applied. The no form of this command returns the policer-control-policy's parent policer maximum rate to max.
Default	no max-rate
Parameters	<i>kilobits-per-second</i> — Defining a kilobits-per-second value is mutually exclusive with the max parameter. The kilobits-per-second value must be defined as an integer that represents the number of kilobytes that the parent policer will be decremented per second. The actual decrement is performed per packet, based on the time that has elapsed since the last packet associated with the parent policer. Values max or 1 to 2000000000 <i>max</i> — The max parameter is mutually exclusive with defining a kilobits-per-second value. When max is specified, the parent policer does not enforce a maximum rate on the aggregate throughput of the child policers. This is the default setting when the policer-control-policy is first created and is the value that the parent policer returns to when no max-rate is executed. In order for the parent policer to be effective, a kilobits-per-second value should be specified.

profile-preferred

Syntax	profile-preferred no profile-preferred
Context	config>qos>plcr-ctrl-plcy>root
Description	The profile-preferred option ensures that the root policer provides a preference to consume its PIR bucket tokens at a given priority level to packets that have their profile state set to in-profile by the output of the child policer CIR bucket.
Default	no profile-preferred

priority-mbs-thresholds

Syntax	priority-mbs-thresholds
Context	config>qos>plcr-ctrl-plcy>root
Description	The priority-mbs-thresholds command contains the root arbiter parent policer's min-thresh-separation command and each priority level's mbs-contribution command that is used to internally derive each priority level's shared-portion and fair-portion values. The system uses each priority level's shared-portion and fair-portion value to calculate each priority level's discard-unfair and discard-all MBS thresholds that enforce priority-sensitive rate-based discards within the root arbiter's parent policer. The priority-mbs-thresholds CLI node always exists and does not need to be created.
Default	None.

min-thresh-separation

Syntax	min-thresh-separation <i>size</i> [bytes kilobytes] no min-thresh-separation
Context	config>qos>plcr-ctrl-plcy>root>priority-mbs-thresholds
Description	The min-thresh-separation command defines the minimum required separation between each in-use discard threshold maintained for each parent policer context associated with the policer-control-policy. The min-thresh-separation value may be overridden on each SAP or sub-profile to which the policy is applied. The system uses the default or specified min-thresh-separation value in order to determine the minimum separation required between each of the of the parent policer discard thresholds. The system enforces the minimum separation based on the following behavior in two ways. The first is determining the size of the shared-portion for each priority level (when the mbs-contribution command's optional fixed keyword is not specified):

- When a parent policer instance's priority level has less than two child policers associated, the shared-portion for the level will be zero.
- When a parent policer instance's priority level has two or more child policers associated, the shared-portion for the level will be equal to the current value of **min-thresh-separation**.

The second function the system uses the **min-thresh-separation** value for is determining the value per priority level for the fair-portion:

- When a parent policer instance's priority level has no child policers associated, the fair-portion for the level will be zero.
- When a parent policer instance's priority level has one child policer associated, the fair-portion will be equal to the maximum of the min-thresh-separation value and the priority level's mbs-contribution value.
- When a parent policer instance's priority level has two or more child policers associated, the fair-portion will be equal to the maximum of the following:
 - **min-thresh-separation** value
 - **mbs-contribution** value less **min-thresh-separation** value

When the **mbs-contribution** command's optional fixed keyword is defined for a priority level within the policy, the system will treat the defined **mbs-contribution** value as an explicit definition of the priority level's MBS. While the system will continue to track child policer associations with the parent policer priority levels, the association counters will have no effect. Instead, the following rules will be used to determine a fixed priority level's shared-portion and fair-portion:

- If a fixed priority level's **mbs-contribution** value is set to zero, both the shared-portion and fair-portion will be set to zero
- If the **mbs-contribution** value is not set to zero:
 - The shared-portion will be set to the current **min-thresh-separation** value
 - The fair-portion will be set to the maximum of the following:
 - min-thresh-separation** value
 - mbs-contribution** value less **min-thresh-separation** value

Each time the **min-thresh-separation** value is modified, the thresholds for all instances of the parent policer created through association with this **policer-control-policy** are reevaluated

Determining the Correct Value for the Minimum Threshold Separation Value

The minimum value for **min-thresh-separation** should be set equal to the maximum size packet that will be handled by the parent policer. This ensures that when a lower priority packet is incrementing the bucket, the size of the increment will not cause the bucket's depth to equal or exceed a higher priority threshold. It also ensures that an unfair packet within a priority level cannot cause the PIR bucket to increment to the discard-all threshold within the priority.

When evaluating maximum packet size, each child policer's per-packet-offset setting should be taken into consideration. If the maximum size packet is 1518 bytes and a per-packet-offset parameter is configured to add 20 bytes per packet, min-thresh-separation should be set to 1538 due to the fact that the parent policer will increment its PIR bucket using the extra 20 bytes.

In most circumstances, a value larger than the maximum packet size is not necessary. Management of priority level aggregate burst tolerance is intended to be implemented using the priority level **mbs-contribution** command. Setting a value larger than the maximum packet size will not adversely affect the policer performance, but it may increase the aggregate burst tolerance for each priority level.



Note: A priority level's shared-portion of the parent policer's PIR bucket depth is only necessary to provide some separation between a lower priority's discard-all threshold and this priority's discard-unfair threshold. It is expected that the burst tolerance for the unfair packets is relatively minimal since the child policers feeding the parent policer priority level all have some amount of fair burst before entering into an FIR exceed or unfair state. The fair burst amount for a priority level is defined using the mbs-contribution command.

The **no** form of this command returns the policy's **min-thresh-separation** value to the default value.

Default no min-thresh-separation

Parameters	<i>size</i> [bytes kilobytes] — The size parameter is required when executing the min-thresh-separation command. It is expressed as an integer and specifies the shared portion in bytes or kilobytes that is selected by the trailing bytes or kilobytes keywords. If both bytes and kilobytes are missing, kilobytes is the assumed value. Setting this value has no effect on parent policer instances where the min-thresh-separation value has been overridden.

Values 0 to 4194304 or **default** (applies to the 7450 ESS)
0 to 16777216 or **default** (applies to the 7750 SR or 7950 XRS)

Default	1536

[bytes | kilobytes] — The **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in kilobytes.

Values **bytes or kilobytes**

Default **kilobytes**

priority

Syntax	priority <i>level</i>
Context	config>qos>plcr-ctrl-plcy>root>priority-mbs-thresholds
Description	The priority level command contains the mbs-contribution configuration command for a given strict priority level. Eight levels are supported numbered 1 through 8 with 8 being the highest strict priority. Each of the eight priority CLI nodes always exists and do not need to be created. While parameters exist for each priority level, the parameters are only applied when the priority level within a parent policer instance is currently supporting child policers.
Default	None.

mbs-contribution

Syntax	mbs-contribution <i>size</i> [bytes kilobytes] [fixed] no mbs-contribution
Context	config>qos>plcr-ctrl-plcy>root>priority-mbs-thresholds>priority
Description	The mbs-contribution command is used to configure the policy-based burst tolerance for a parent policer instance created when the policy is applied to a SAP, or a subscriber context for the 7450 ESS and 7750 SR, or a 7950 XRS multiservice site. The system uses the parent policer's min-thresh-separation value, the priority level's mbs-contribution value, and the number of child policers currently attached to the priority level to derive the priority level's shared-portion and fair-portion of burst tolerance within the local priority level. The shared-portion and fair-portions for each priority level are then used by the system to calculate each priority level's discard-unfair threshold and discard-all threshold. The mbs-contribution is the minimum separation between two adjacent active discard-all thresholds.

The value for a priority level's **mbs-contribution** within the policer-control-policy may be overridden on the SAP, or subscriber sub-profile (applies to the 7450 ESS and 7750 SR) or multiservice site (for 7950 XRS) where the policy is applied in order to allow fine tuning of the discard-unfair and discard-all thresholds relevant to the needs of the local child policers on the object.

Accumulative Nature of Burst Tolerance for a Parent Policer Priority Level

When defining **mbs-contribution**, the specified size may only be a portion of the burst tolerance associated with the priority level. The packets associated with the priority level share the burst tolerance of lower within the parent policer. As the parent policer PIR bucket depth increases during congestion, the lower priority packets eventually experience discard based on each priority's discard-unfair and discard-all thresholds. Assuming congestion continues when all the lower priority packets have been prevented from consuming bucket depth, the burst tolerance for the priority level will be consumed by its own packets and any packets associated with higher priorities.

The Effect of Fair and Unfair Child Policer Traffic at a Parent Policer Priority Level

The system continually monitors the offered rate of each child policer on each parent policer priority level and detects when the policer is in a congested state (the aggregate offered load is greater than the decrement rate defined on the parent policer). As previously stated, the result of congestion is that the parent policer's bucket depth will increase until it eventually hovers around either a discard-unfair or discard-all threshold belonging to one of the priority levels. This threshold is the point where enough packets are being discarded that the increment rate and decrement rate begin to even out. If only a single child policer is associated with the priority level, the discard-unfair threshold is not used since fairness is only applicable when multiple child policers are competing at the same priority level.

When multiple child policers are sharing the congested priority level, the system uses the offered rates and the parenting parameters of each child to determine the fair rate per child when the parent policer is unable to meet the bandwidth needs of each child. The fair rate represents the amount of bandwidth that each child at the priority level should receive relative to the other children at the same level according to the policer control policy instance managing the child policers. This fair rate is applied as the decrement rate for each child's FIR bucket. Changing a child's FIR rate does not modify the amount of packets forwarded by the parent policer for the child's priority level. It just modifies the forwarded ratio between the children on that priority level. Since each child FIR bucket has some level of burst tolerance before marking its packets as unfair, the current parent policer bucket depth may at times rise above the discard-unfair threshold. The mbs-contribution value provides a means to define how much separation is provided between the priority level's discard-unfair and discard-all threshold to allow the parent policer to absorb some amount of FIR burst before reaching the priority's discard-all threshold.

This level of fair aggregate burst tolerance is based on the decrement rate of the parent policer's PIR bucket while the individual fair bursts making up the aggregate are based on each child's FIR decrement rate. The aggregate fair rate of the priority level is managed by the system with consideration of the current rate of traffic in higher priority levels. In essence, the system ensures that for each iteration of the child FIR rate calculation, the sum of the child FIR decrement rates plus the sum of the higher priority traffic increment rates equals the parent policers decrement rate. This means that dynamic amounts of higher priority traffic can be ignored when sizing a lower priority's fair aggregate burst tolerance. Consider the following:

- The parent policer decrement rate is set to 20 Mb/s (max-rate 20,000).
- A priority level's fair burst size is set to 30 kbytes (mbs-contribution 30 kbytes).
- Higher priority traffic is currently taking 12 Mb/s.
- The priority level has three child policers attached.
- Each child's PIR MBS is set to 10 kbytes, which makes each child's FIR MBS 10 kbytes.
- The children want 10 Mb/s, but only 8 Mb/s is available
- Based on weights, the children's FIR rates are set as follows:

	FIR Rate	FIR MBS
Child 1	4 Mb/s	10 kbytes
Child 2	3 Mb/s	10 kbytes

	FIR Rate	FIR MBS
Child 3	1 Mb/s	10 kbytes

The 12 Mb/s of the higher priority traffic and the 8 Mb/s of fair traffic equal the 20 Mb/s decrement rate of the parent policer.

It is clear that the higher priority traffic is consuming 12 Mb/s of the parent policer's decrement rate, leaving 8 Mb/s of decrement rate for the lower priority's fair traffic.

- The burst tolerance of child 1 is based on 10 kbytes above 4 Mb/s.
- The burst tolerance of child 2 is based on 10 kbytes above 3 Mb/s.
- The burst tolerance of child 3 is based on 10 kbytes above 1 Mb/s.

If all three children burst simultaneously (unlikely), they will consume 30 kbytes above 8 Mb/s. This is the same as the remaining decrement rate after the higher priority traffic.

Parent Policer Total Burst Tolerance and Downstream Buffering

The highest in-use priority level's discard-all threshold is the total burst tolerance of the parent policer. In some cases, the parent policer represents downstream bandwidth capacity and the max-rate of the parent policer is set to prevent overrunning the downstream bandwidth. The burst tolerance of the parent policer defines how much more traffic may be sent beyond the downstream scheduling capacity. In the worst case scenario, when the downstream buffering is insufficient to handle the total possible burst from the parent policer, downstream discards based on lack of buffering may occur. However, in all likelihood, this is not the case.

In most cases, lower priority traffic in the policer will be responsible for the greater part of congestion above the parent policer rate. Since this traffic is discarded with a lower threshold, this lowers the effective burst tolerance even while the highest priority traffic is present.

Configuring a Priority Level's MBS Contribution Value

In the most conservative case, a priority level's **mbs-contribution** value may be set to be greater than the sum of child policer's mbs and one max-size-frame per child policer. This ensures that even in the absolute worst case where all the lower priority levels are simultaneously bursting to the maximum capacity of each child, enough burst tolerance for the priority's children will exist if they also burst to their maximum capacity.

Since simply adding up all the child policer's PIR MBS values may result in large overall burst tolerances that are not ever likely to be needed, consider some level of burst oversubscription when configuring the **mbs-contribution** value for each priority level. The amount of oversubscription should be determined based on the needs of each priority level.

Using the Fixed Keyword to Create Deterministic Parent Policer Discard Thresholds

In the default behavior, the system ignores the **mbs-contribution** values for a priority level on a subscriber for 7450 ESS and 7750 SR, or a multiservice site for 7950 XRS, or SAP parent policer when a child policer is not currently associated with the level. This prevents additional burst tolerance from being added to higher priority traffic within the parent policer.

This does cause fluctuations in the defined threshold values when child policers are added or removed from a parent policer instance. If this behavior is undesirable, the fixed keyword may be used that causes the **mbs-contribution** value to always be included in the calculation of parent policer's discard thresholds. The defined **mbs-contribution** value may be overridden on a subscriber sla-profile for the 7450 ESS and 7750 SR, or on a multiservice site for the 7950 XRS, or on a SAP instance, but the fixed nature of the contribution cannot be overridden.

If the defined **mbs-contribution** value for the priority level is zero, the priority level will have no effect on the parent policer's defined discard thresholds. A packet associated with the priority level will use the next lower priority level's discard-unfair and discard-all thresholds.

The **no** form of this command returns the policy's priority level's MBS contribution to the default value. When changed, the thresholds for the priority level and all higher priority levels for all instances of the parent policer will be recalculated.

Default no mbs-contribution

Parameters *size [bytes | kilobytes]* — The size parameter is required when executing the **mbs-contribution** command. It is expressed as an integer and specifies the priority's specific portion amount of accumulative MBS for the priority level in bytes or kilobytes that is selected by the trailing **bytes** or **kilobytes** keywords. If both **bytes** and **kilobytes** are missing, **kilobytes** is assumed. Setting this value has no effect on parent policer instances where the priority level's **mbs-contribution** value has been overridden.

Values 0 to 4194304 or **default** (applies to the 7450 ESS)
0 to 16777216 or **default** (applies to the 7750 SR and 7950 XRS)

Default 8 kilobytes

bytes | kilobytes: — The **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. When specified, size is interpreted as specifying the size of min-thresh-separation in kilobytes.

Default kilobytes

fixed — The optional fixed keyword is used to force the inclusion of the defined **mbs-contribution** value in the parent policer's discard threshold calculations. If the **mbs-contribution** command is executed without the **fixed** keyword, the fixed calculation behavior for the priority level is removed.

tier

Syntax tier {1 | 2}

Context config>qos>policer-control-policy

Description	This command is used to create, configure, and delete tiered arbiters. Two tiers are supported that always exist, specified as tier 1 and tier 2. Tiered arbiters enable the creation of a bandwidth control hierarchy for managing child policers in an arbitrary fashion. Each arbiter enables parenting of child policers within eight strict levels of priority and a maximum aggregate rate may be defined for the children that the arbiter will enforce. Arbiters created on tier 1 are automatically parented to the root arbiter that is always present. Arbiters created on tier 2 default to the root arbiter as parent, but can also be explicitly parented to a tier 2 arbiter. Child policers associated with an instance of the policer-control-policy can be parented to any tiered arbiter or to the root arbiter.
Default	None.

arbiter

Syntax	arbiter <i>arbiter-name</i> [create] no arbiter <i>arbiter-name</i>
Context	config>qos>plcr-ctrl-plcy>tier
Description	This command is used to create an arbiter within the context of tier 1 or tier 2. An arbiter is a child policer bandwidth control object that manages the throughput of a set of child policers. An arbiter allows child policers or other arbiters to parent to one of eight strict levels. Each arbiter is itself parented to either another tiered arbiter or to the root arbiter. The root arbiter starts with its defined maximum rate and distributes the bandwidth to its directly attached child policers and arbiters beginning with priority 8. As the children at each priority level are distributed bandwidth according to their needs and limits, the root proceeds to the next lower priority until either all children's needs are met or it runs out of bandwidth. The bandwidth given to a tiered arbiter is then divided between that arbiter's children (child policers or a tier 2 arbiter) in the same fashion. A tiered arbiter may also have a rate limit defined that limits the amount of bandwidth it may receive from its parent. An arbiter that is currently parented by another arbiter cannot be deleted. Each time the policer-control-policy is applied to either a SAP, or a subscriber (through association with a sub-profile that has the policy applied), or a multiservice site, an instance of the parent policer and the arbiters is created. Any child policer that uses the arbiter's name in its parenting command will be associated with the arbiter instance. The child policer will also become associated with any arbiter to which its parent arbiter is parented (grandparent). Having child policers parented to an arbiter does not prevent that arbiter from being removed from the policer-control-policy. When removed, the child policers become orphaned. You can create up to 31 tiered arbiters within the policer-control-policy on either tier 1 or tier 2 (in addition to the arbiter).

The **no** form of this command is used to remove an arbiter from tier 1 or tier 2. If the specified arbiter does not exist, the command returns without an error. If the specified arbiter is currently specified as the parent for another arbiter, the command will fail. When an arbiter is removed from a **policer-control-policy**, all instances of the arbiter will also be removed. Any child policers currently parented to the arbiter instance will become orphans and will not be bandwidth managed by the policer control policy instances parent policer.

Default None.

Parameters *arbiter-name* — Any unique name within the policy. Up to 31 arbiters may be created.

rate

Syntax **rate** {*kilobits-per-second* | **max**}

Context config>qos>plcr-ctrl-plcy>tier>arbiter

Description This command is used to define the maximum bandwidth an instance of the arbiter can receive from its parent tier 1 arbiter or the root arbiter. The arbiter instance enforces this limit by calculating the bandwidth each of its child policers should receive relative to their offered loads, parenting parameters, and individual rate limits, and using that derived rate as a child PIR decrement rate override. The override will not exceed the child policer's administrative rate limit and the aggregate of all the child PIR decrement rates will not exceed the specified arbiter rate limit.

The arbiter's policy defined rate value may be overridden at the SAP or sub-profile where the **policer-control-policy** is applied. Specifying an override prevents the arbiter from being removed from the policer control policy until the override is removed.

The **no** form of this command is used to remove a rate limit from the arbiter at the policer control policy level. The policy level rate limit for the arbiter will return to the default value of **max**. The **no rate** command has no effect on instances of the arbiter where a rate limit override has been defined.

Default rate max

Parameters *kilobits-per-second* — **max** or 1 to 2000000000

The *kilobits-per-second* parameter is mutually exclusive with the **max** keyword. When specifying a value for *kilobits-per-second*, enter an integer representing the rate limit in kilobits per second.

max — The **max** keyword is mutually exclusive with the *kilobits-per-second* parameter. When **max** is specified, the arbiter does not enforce a rate limit on its child policers or arbiters other than the individual rate limits enforced at the child level.

parent

Syntax **parent** {**root** | *arbiter-name*} [**level** *priority-level*] [**weight** *weight-within-level*]

	no parent
Context	config>qos>plcr-ctrl-plcy>tier>arbiter
Description	This command is used to define from where the tiered arbiter receives bandwidth. Both tier 1 and tier 2 arbiters default to parenting to the root arbiter. Tier 2 arbiters may be modified to parent to a tier 1 arbiter. The tier 1 arbiter parent cannot be changed. The parent command is also used to define the parenting parameters. Each child arbiter attaches to its parent on one of the parent's eight strict levels. Level 1 is the lowest and 8 is the highest. The level attribute is used to define which level the child arbiter uses on its parent. The parent distributes its available bandwidth based on strict priority starting with priority level 8 and proceeding towards level 1. The weight attribute is used to define how multiple children at the same parent strict level compete when insufficient bandwidth exists on the parent for that level. Each child's weight is divided by the sum of the active children's weights and the result is multiplied by the available bandwidth. If a child cannot receive its entire weighted fair share of bandwidth due to a defined child rate limit, the remainder of its bandwidth is distributed between the other children based on their weights. The no form of this command is used to return the tiered arbiter to the default parenting behavior.
Default	parent "root" level 1 weight 1
Parameters	root — The root keyword is mutually exclusive with the <i>arbiter-name</i> parameter. In tier 1, <i>arbiter-name</i> is not allowed and only root is accepted. When root is specified, the arbiter will receive all bandwidth directly from the root arbiter. This is the default parent for tiered arbiters. <i>arbiter-name</i> — The <i>arbiter-name</i> parameter is mutually exclusive with the root keyword. In tier 1, <i>arbiter-name</i> is not allowed and only root is accepted. The specified <i>arbiter-name</i> must exist within the policer-control-policy at tier 1 or the parent command will fail. When a tiered arbiter is acting as a parent for another tiered arbiter, the parent arbiter cannot be removed from the policy. The child arbiter will receive all bandwidth directly from its parent arbiter (that receives bandwidth from the root arbiter). level <i>priority-level</i> — The level <i>priority-level</i> keyword and parameter are optional when executing the parent command. When level is not specified, a default level of 1 is used in the parent arbiter. When level is specified, the <i>priority-level</i> parameter must be specified as an integer value from 1 through 8. weight <i>weight-within-level</i> — The weight <i>weight-within-level</i> 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 <i>weight-within-level</i> parameter must be specified as an integer value from 1 through 100.

post-policer-mapping

Syntax	post-policer-mapping <i>mapping-policy-name</i> [create] no post-policer-mapping <i>mapping-policy-name</i>
Context	config>qos
Description	This command configures a post-policer mapping policy which is used to remap a packet's forwarding class and profile state to another forwarding class and profile state for post-policer traffic. A post-policer mapping policy is created without any forwarding class or profile remapping statements. If an empty policy is applied to a SAP-egress QoS policy, no remapping occurs. The no form of this command deletes the post-policer mapping policy. A post-policer mapping policy can only be deleted if there are no references to it.
Parameters	<i>mapping-policy-name</i> — Specifies the name of the post-policer mapping policy, up to 32 characters.

fc

Syntax	fc <i>fc-name</i> profile <i>profile</i> [create] no fc <i>fc-name</i> profile <i>profile</i>
Context	config>qos>post-policer-mapping
Description	This command specifies the forwarding class and profile state of an egress policed packet that is to be mapped to another forwarding class and profile, where the profile state is that of the resulting profile after the packet has been processed by the egress policer. The new forwarding class and profile state is configured using the maps-to command. The traffic remarking is based on the marking configured for the forwarding class and profile of the traffic after being policed but before it is remapped. The no form of this command deletes the forwarding class and profile remapping statement, including the maps-to command.
Parameters	<i>fc-name</i> — Specifies one of the eight forwarding classes supported by the system. Values be, l2, af, l1, h2, ef, h1, nc <i>profile</i> — Specifies one of the egress packet profile states. Values exceed, in, inplus, out

maps-to

Syntax	maps-to fc <i>fc-name</i> profile <i>profile</i>
Context	config>qos>post-policer-mapping>fc
Description	This command remaps the forwarding class and profile state of an egress policed packet that is to be mapped to another forwarding class and profile, where the profile state is that of the resulting profile after the packet has been processed by the egress policer. The new forwarding class is used to select the egress queue on which the post-policer traffic is placed. The new profile is used to determine the congestion control handling in that queue, specifically the drop tail or slope that is applied to the traffic. The maps-to command parameters can be overwritten by reissuing the command with a different FC or profile. The traffic remarking is based on the marking configured for the forwarding class and profile of the traffic after being policed but before it is remapped.
Parameters	<i>fc-name</i> — Specifies one of the eight forwarding classes supported by the system. Values be, l2, af, l1, h2, ef, h1, nc <i>profile</i> — Specifies one of the egress packet profile states. Values exceed, in, inplus, out

17.13.2.1.3 Operational Commands

post-policer-mapping

Syntax	post-policer-mapping <i>src-name</i> <i>dst-name</i> [overwrite]
Context	config>qos>copy
Description	This command copies an existing post-policer mapping policy to another policy name. 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	<i>src-name</i> — Specifies the source policy name that the copy command will attempt to copy from. <i>dst-name</i> — Specifies the destination policy name to which the command will copy a duplicate of the policy.

overwrite — Specifies that the existing destination policy is to be replaced. 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 name exists.

17.13.2.1.4 Show Commands

post-policer-mapping

Syntax	post-policer-mapping [<i>post-policer-name</i>] [association detail]
Context	show>qos
Description	This command displays information about post-policer mapping policies. Omitting the policy name will list all policies.
Parameters	<i>post-policer-name</i> — Specifies the name of the post-policer mapping policy, up to 32 characters. association — Displays the entities associated with the post-policer mapping policy. detail — Displays the detailed information for the post-policer mapping policy.
Output	This section shows sample output for post-policer mapping.

Sample Output

```
*A:PE# show qos post-policer-mapping
=====
Post Policer Mapping Summary
=====
Policy                               Description
-----
ppml                                Map FC EF/exceed to BE/out
-----
Policies
=====
*A:PE#

*A:PE# show qos post-policer-mapping detail
=====
Post Policer Mapping Information
=====
Policy Name      : ppml
Description      : Map FC EF/exceed to BE/out
-----
Post Policer Mapping Information
-----
FC               Profile           Mapped-FC           Mapped-Profile
-----
ef               Exceed             be                   Out
-----
```

```
-----  
Sap Egress Associations  
-----  
Policy ID  
-----  
10  
-----  
=====
```

18 Frequently Used QoS Terms

18.1 Overview

This section provides definitions for frequently used QoS terminology.

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.

18.2 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.

18.3 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.

18.4 CBS

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

18.5 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.

18.6 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.

18.7 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, 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.

18.8 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.

18.9 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.

18.10 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 XMA or MDA forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.

18.11 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.

18.12 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.

18.13 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) than the child scheduler. Parent associations for queues are only checked when, when an instance of the queue is created on a SAP.

18.14 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.

18.15 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 (kb/s).

- 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.

18.16 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.

18.17 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 multiservice site or to a service Service Access Point (SAP). When 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.

18.18 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.

18.19 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.

18.20 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 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 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.

18.21 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.

19 Standards and Protocol Support



Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

Bidirectional Forwarding Detection (BFD)

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

Border Gateway Protocol (BGP)

draft-hares-idr-update-attr-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*

draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*

draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*

draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*

draft-ietf-idr-flowspec-interfaceset-03, *Applying BGP flowspec rules on a specific interface set*

draft-ietf-idr-flowspec-path-redirect-05, *Flowspec Indirection-id Redirect (localised ID)*

draft-ietf-idr-flowspec-redirect-ip-02, *BGP Flow-Spec Redirect to IP Action*

draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*

draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*

draft-uttaro-idr-bgp-persistence-03, *Support for Long-lived BGP Graceful Restart*

RFC 1772, *Application of the Border Gateway Protocol in the Internet*

RFC 1997, *BGP Communities Attribute*

RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*

RFC 2439, *BGP Route Flap Damping*

RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*

RFC 2858, *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 BGP-4*

RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

RFC 4360, *BGP Extended Communities Attribute*

RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*

RFC 4486, *Subcodes for BGP Cease Notification Message*
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 (helper mode)*
RFC 4760, *Multiprotocol Extensions for BGP-4*
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, *Autonomous System Confederations for BGP*
RFC 5291, *Outbound Route Filtering Capability for BGP-4*
RFC 5396, *Textual Representation of Autonomous System (AS) Numbers (asplain)*
RFC 5549, *Advertising IPv4 Network Layer Reachability Information with an IPv6
Next Hop*
RFC 5575, *Dissemination of Flow Specification Rules*
RFC 5668, *4-Octet AS Specific BGP Extended Community*
RFC 6286, *Autonomous-System-Wide Unique BGP Identifier for BGP-4*
RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*
RFC 6811, *Prefix Origin Validation*
RFC 6996, *Autonomous System (AS) Reservation for Private Use*
RFC 7311, *The Accumulated IGP Metric Attribute for BGP*
RFC 7607, *Codification of AS 0 Processing*
RFC 7674, *Clarification of the Flowspec Redirect Extended Community*
RFC 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE)
Information Using BGP*
RFC 7854, *BGP Monitoring Protocol (BMP)*
RFC 7911, *Advertisement of Multiple Paths in BGP*
RFC 7999, *BLACKHOLE Community*
RFC 8092, *BGP Large Communities Attribute*

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)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1ag, *Connectivity Fault Management*

IEEE 802.1ah, *Provider Backbone Bridges*

IEEE 802.1ak, *Multiple Registration Protocol*

IEEE 802.1aq, *Shortest Path Bridging*

IEEE 802.1ax, *Link Aggregation*

IEEE 802.1D, *MAC Bridges*

IEEE 802.1p, *Traffic Class Expediting*

IEEE 802.1Q, *Virtual LANs*

IEEE 802.1s, *Multiple Spanning Trees*

IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*

IEEE 802.1X, *Port Based Network Access Control*

IEEE 802.3ab, *1000BASE-T*

IEEE 802.3ac, *VLAN Tag*

IEEE 802.3ad, *Link Aggregation*

IEEE 802.3ae, *10 Gb/s Ethernet*

IEEE 802.3ah, *Ethernet in the First Mile*

IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*

IEEE 802.3i, *Ethernet*

IEEE 802.3u, *Fast Ethernet*

IEEE 802.3x, *Ethernet Flow Control*

IEEE 802.3z, *Gigabit Ethernet*

ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*

ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*

ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

Ethernet VPN (EVPN)

draft-ietf-bess-evpn-ac-df-01, *AC-Influenced Designated Forwarder Election for EVPN*

draft-ietf-bess-evpn-pref-df-01, *Preference-based EVPN DF Election*

draft-ietf-bess-evpn-prefix-advertisement-11, *IP Prefix Advertisement in EVPN*

draft-ietf-bess-evpn-proxy-arp-nd-04, Operational Aspects of Proxy-ARP/ND in EVPN Networks
draft-ietf-bess-evpn-vpls-seamless-integ-03, (PBB-)EVPN Seamless Integration with (PBB-)VPLS
draft-snr-bess-pbb-evpn-isid-cmacflush-01, PBB-EVPN ISID-based CMAC-Flush
RFC 7432, BGP MPLS-Based Ethernet VPN
RFC 7623, Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)
RFC 8214, Virtual Private Wire Service Support in Ethernet VPN
RFC 8317, Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) an Provider Backbone Bridging EVPN (PBB-EVPN)
RFC 8365, A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)

Frame Relay

ANSI T1.617 Annex D, DSS1 - Signalling Specification For Frame Relay Bearer Service
FRF.1.2, PVC User-to-Network Interface (UNI) Implementation Agreement
FRF.12, Frame Relay Fragmentation Implementation Agreement
FRF.16.1, Multilink Frame Relay UNI/NNI Implementation Agreement
FRF.5, Frame Relay/ATM PVC Network Interworking Implementation
FRF2.2, PVC Network-to-Network Interface (NNI) Implementation Agreement
ITU-T Q.933 Annex A, Additional procedures for Permanent Virtual Connection (PVC) status management

Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information
RFC 3471, Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description
RFC 3473, Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions
RFC 4204, Link Management Protocol (LMP)
RFC 4208, Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model
RFC 4872, RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery
RFC 5063, Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart (helper mode)

gRPC Remote Procedure Calls (gRPC)

gnmi.proto, *gRPC Network Management Interface (gNMI), version 0.4.0*

*gRPC Network Management Interface (gNMI), Capabilities, Get, Set, Subscribe
(ONCE, SAMPLE, ON_CHANGE)*

Intermediate System to Intermediate System (IS-IS)

draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to
Intermediate system intra-domain routing information exchange protocol for
use in conjunction with the protocol for providing the connectionless-mode
Network Service (ISO 8473)*

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 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*

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 (helper mode)*

RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label
Switching (GMPLS)*

RFC 5308, *Routing IPv6 with IS-IS*

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 6232, *Purge Originator Identification TLV for IS-IS*

RFC 6233, *IS-IS Registry Extension for Purges*

RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*
RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*
RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability*
RFC 8202, *IS-IS Multi-Instance* (single topology)

Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*
RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
RFC 7431, *Multicast-Only Fast Reroute*
RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*
RFC 768, *User Datagram Protocol*
RFC 793, *Transmission Control Protocol*
RFC 854, *Telnet Protocol Specifications*
RFC 1350, *The TFTP Protocol (revision 2)*
RFC 2347, *TFTP Option Extension*
RFC 2348, *TFTP Blocksize Option*
RFC 2349, *TFTP Timeout Interval and Transfer Size Options*
RFC 2428, *FTP Extensions for IPv6 and NATs*
RFC 2784, *Generic Routing Encapsulation (GRE)*
RFC 2890, *Key and Sequence Number Extensions to GRE*
RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*
RFC 4251, *The Secure Shell (SSH) Protocol Architecture*
RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)
RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*
RFC 4254, *The Secure Shell (SSH) Connection Protocol*
RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*
RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*
RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer (ECDSA)*
RFC 5925, *The TCP Authentication Option*
RFC 5926, *Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)*
RFC 6398, *IP Router Alert Considerations and Usage (MLD)*

RFC 6528, *Defending against Sequence Number Attacks*

Internet Protocol (IP) — Multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* (version 1)

draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*

draft-ietf-bier-mvpn-11, *Multicast VPN Using BIER*

draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*

draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*

draft-ietf-mboned-mtrace-v2-17, *Mtrace Version 2: Traceroute Facility for IP Multicast*

RFC 1112, *Host Extensions for IP Multicasting*

RFC 2236, *Internet Group Management Protocol, Version 2*

RFC 2365, *Administratively Scoped IP Multicast*

RFC 2375, *IPv6 Multicast Address Assignments*

RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*

RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*

RFC 3376, *Internet Group Management Protocol, Version 3*

RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*

RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*

RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*

RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*

RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)* (auto-RP groups)

RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*

RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (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)*

RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*

RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*

RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*

RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*

RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*

RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6513, *Multicast in MPLS/BGP IP VPNs*

RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*

RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*

RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*

RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*

RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*

RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*

RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*

RFC 7761, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*

RFC 8279, *Multicast Using Bit Index Explicit Replication (BIER)*

RFC 8296, *Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks (MPLS encapsulation)*

RFC 8401, *Bit Index Explicit Replication (BIER) Support via IS-IS*

Internet Protocol (IP) — Version 4

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1191, *Path MTU Discovery (router specification)*

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*
RFC 1534, *Interoperation between DHCP and BOOTP*
RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*
RFC 1812, *Requirements for IPv4 Routers*
RFC 1918, *Address Allocation for Private Internets*
RFC 2003, *IP Encapsulation within IP*
RFC 2131, *Dynamic Host Configuration Protocol*
RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
RFC 2401, *Security Architecture for Internet Protocol*
RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*
RFC 3046, *DHCP Relay Agent Information Option (Option 82)*
RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*
RFC 4884, *Extended ICMP to Support Multi-Part Messages (ICMPv4 and ICMPv6 Time Exceeded)*

Internet Protocol (IP) — Version 6

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*
RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*
RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*
RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3587, *IPv6 Global Unicast Address Format*
RFC 3596, *DNS Extensions to Support IP version 6*
RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*
RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*
RFC 3971, *SEcure Neighbor Discovery (SEND)*
RFC 3972, *Cryptographically Generated Addresses (CGA)*
RFC 4007, *IPv6 Scoped Address Architecture*
RFC 4193, *Unique Local IPv6 Unicast Addresses*
RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*
RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*
RFC 4862, *IPv6 Stateless Address Autoconfiguration (router functions)*

RFC 4890, *Recommendations for Filtering ICMPv6 Messages in Firewalls*
RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*
RFC 5007, *DHCPv6 Leasequery*
RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*
RFC 5722, *Handling of Overlapping IPv6 Fragments*
RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6 (IPv6)*
RFC 5952, *A Recommendation for IPv6 Address Text Representation*
RFC 6092, *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service* (Internet Control and Management, Upper-Layer Transport Protocols, UDP Filters, IPsec and Internet Key Exchange (IKE), TCP Filters)
RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*
RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*
RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*
RFC 8201, *Path MTU Discovery for IP version 6*

Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*
draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*
RFC 2401, *Security Architecture for the Internet Protocol*
RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*
RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*
RFC 2406, *IP Encapsulating Security Payload (ESP)*
RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*
RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*
RFC 2409, *The Internet Key Exchange (IKE)*
RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*
RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*
RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*
RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*
RFC 3947, *Negotiation of NAT-Traversal in the IKE*

RFC 3948, *UDP Encapsulation of IPsec ESP Packets*
RFC 4106, *The Use of Galois/Counter Mode (GCM) in IPsec ESP*
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 4301, *Security Architecture for the Internet Protocol*
RFC 4303, *IP Encapsulating Security Payload*
RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*
RFC 4308, *Cryptographic Suites for IPsec*
RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*
RFC 4543, *The Use of Galois Message Authentication Code (GMAC) in IPsec ESP and AH*
RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPSec*
RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*
RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*
RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*
RFC 5282, *Using Authenticated Encryption Algorithms with the Encrypted Payload of the IKEv2 Protocol*
RFC 5903, *ECP Groups for IKE and IKEv2*
RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*
RFC 6379, *Suite B Cryptographic Suites for IPsec*
RFC 6380, *Suite B Profile for Internet Protocol Security (IPsec)*
RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*
RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*
RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*
RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*
RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2 (IKEv2)*
RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

Label Distribution Protocol (LDP)

draft-ietf-mpls-ldp-ip-pw-capability-09, Controlling State Advertisements Of Non-negotiated LDP Applications

draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities

draft-pdutta-mpls-ldp-v2-00, LDP Version 2

draft-pdutta-mpls-mldp-up-redundancy-00, Upstream LSR Redundancy for Multipoint LDP Tunnels

draft-pdutta-mpls-multi-ldp-instance-00, Multiple LDP Instances

draft-pdutta-mpls-tldp-hello-reduce-04, Targeted LDP Hello Reduction

RFC 3037, LDP Applicability

RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (helper mode)

RFC 5036, LDP Specification

RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs)

RFC 5443, LDP IGP Synchronization

RFC 5561, LDP Capabilities

RFC 5919, Signaling LDP Label Advertisement Completion

RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

RFC 6512, Using Multipoint LDP When the Backbone Has No Route to the Root

RFC 6826, Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

RFC 7032, LDP Downstream-on-Demand in Seamless MPLS

RFC 7552, Updates to LDP for IPv6

Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions

RFC 2661, Layer Two Tunneling Protocol "L2TP"

RFC 2809, Implementation of L2TP Compulsory Tunneling via RADIUS

RFC 3438, Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update

RFC 3931, Layer Two Tunneling Protocol - Version 3 (L2TPv3)

RFC 4719, Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)

RFC 4951, Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"

Management

draft-ietf-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-vrrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6 (IPv6)*

ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*

ianagmplstc-mib, *IANA-GMPLS-TC-MIB*

ianaiftype-mib, *IANAifType-MIB*

ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*

IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*

IEEE8021-PAE-MIB, *IEEE 802.1X MIB*

IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*

LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*

openconfig-bgp.yang version 3.0.1, *BGP Module*

openconfig-bgp-common.yang version 3.0.1, *BGP Common Module*

openconfig-bgp-common-multiprotocol.yang version 3.0.1, *BGP Common Multiprotocol Module*

openconfig-bgp-common-structure.yang version 3.0.1, *BGP Common Structure Module*

openconfig-bgp-global.yang version 3.0.1, *BGP Global Module*

openconfig-bgp-neighbor.yang version 3.0.1, *BGP Neighbor Module*

openconfig-bgp-peer-group.yang version 3.0.1, *BGP Peer Group Module*

openconfig-bgp-policy.yang version 4.0.1, *BGP Policy Module*

openconfig-if-aggregate.yang version 2.0.0, *Interfaces Aggregated Model*

openconfig-if-ethernet.yang version 2.0.0, *Interfaces Ethernet Model*

openconfig-if-ip.yang version 2.0.0, *Interfaces IP Module*

openconfig-if-ip-ext.yang version 2.0.0, *Interfaces IP Extensions Module*

openconfig-interfaces.yang version 2.0.0, *Interfaces Module*

openconfig-isis.yang version 0.3.0, *IS-IS Module*
openconfig-isis-lsp.yang version 0.3.0, *IS-IS LSP Module*
openconfig-isis-routing.yang version 0.3.0, *IS-IS Routing Module*
openconfig-lacp.yang version 1.1.0, *LACP Module*
openconfig-lldp.yang version 0.1.0, *LLDP Module*
openconfig-local-routing.yang version 1.0.1, *Local Routing Module*
openconfig-network-instance.yang version 0.8.0, *Network Instance Module*
openconfig-routing-policy.yang version 3.0.0, *Routing Policy Module*
openconfig-vlan.yang version 2.0.0, *VLAN Module*
RFC 1157, *A Simple Network Management Protocol (SNMP)*
RFC 1212, *Concise MIB Definitions*
RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*
RFC 1215, *A Convention for Defining Traps for use with the SNMP*
RFC 1724, *RIP Version 2 MIB Extension*
RFC 1901, *Introduction to Community-based SNMPv2*
RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*
RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*
RFC 2206, *RSVP Management Information Base using SMIv2*
RFC 2213, *Integrated Services Management Information Base using SMIv2*
RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*
RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*
RFC 2515, *Definitions of Managed Objects for ATM Management*
RFC 2570, *SNMP Version 3 Framework*
RFC 2571, *An Architecture for Describing SNMP Management Frameworks*
RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*
RFC 2573, *SNMP Applications*
RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*
RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*
RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
RFC 2579, *Textual Conventions for SMIv2*
RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*

RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*

RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4)*

RFC 3419, *Textual Conventions for Transport Addresses*

RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*

RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*

RFC 4220, *Traffic Engineering Link Management Information Base*
RFC 4273, *Definitions of Managed Objects for BGP-4*
RFC 4292, *IP Forwarding Table MIB*
RFC 4293, *Management Information Base for the Internet Protocol (IP)*
RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*
RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*
RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms (TLS)*
RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*
RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*
RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*
RFC 5102, *Information Model for IP Flow Information Export*
RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2 (TLS client, RSA public key)*
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*
RFC 6991, *Common YANG Data Types*
RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*
RFC 7950, *The YANG 1.1 Data Modeling Language*
SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

Multiprotocol Label Switching — Transport Profile (MPLS-TP)

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 Transport Profile (MPLS-TP) Identifiers*
RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*
RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*
RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*
RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*
RFC 6478, *Pseudowire Status for Static Pseudowires*

RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

Multiprotocol Label Switching (MPLS)

draft-ietf-teas-sr-rsvp-coexistence-rec-02, *Recommendations for RSVP-TE and Segment Routing LSP co-existence*

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

RFC 7510, *Encapsulating MPLS in UDP*

Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

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 6887, *Port Control Protocol (PCP)*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7915, *IP/ICMP Translation Algorithm*

Network Configuration Protocol (NETCONF)

RFC 5277, *NETCONF Event Notifications*

RFC 6022, *YANG Module for NETCONF Monitoring*
RFC 6241, *Network Configuration Protocol (NETCONF)*
RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*
RFC 6243, *With-defaults Capability for NETCONF*
RFC 7895, *YANG Module Library*

Open Shortest Path First (OSPF)

draft-ietf-ospf-ospfv3-lsa-extend-13, *OSPFv3 LSA Extendibility*
RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*
RFC 1765, *OSPF Database Overflow*
RFC 2328, *OSPF Version 2*
RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*
RFC 3509, *Alternative Implementations of OSPF Area Border Routers*
RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart (helper mode)*
RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*
RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*
RFC 4552, *Authentication/Confidentiality for OSPFv3*
RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 5185, *OSPF Multi-Area Adjacency*
RFC 5187, *OSPFv3 Graceful Restart (helper mode)*
RFC 5243, *OSPF Database Exchange Summary List Optimization*
RFC 5250, *The OSPF Opaque LSA Option*
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
RFC 5340, *OSPF for IPv6*
RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*
RFC 5838, *Support of Address Families in OSPFv3*
RFC 6987, *OSPF Stub Router Advertisement*
RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*
RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

OpenFlow

TS-007, *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*

draft-ietf-pce-segment-routing-08, *PCEP Extensions for Segment Routing*

draft-ietf-pce-stateful-pce-14, *PCEP Extensions for Stateful PCE*

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*

RFC 1661, *The Point-to-Point Protocol (PPP)*

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 2153, *PPP Vendor Extensions*

RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*

RFC 2615, *PPP over SONET/SDH*

RFC 2686, *The Multi-Class Extension to Multi-Link PPP*

RFC 2878, *PPP Bridging Control Protocol (BCP)*

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*

RFC 5072, *IP Version 6 over PPP*

Policy Management and Credit Control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points (Gx support as it applies to wireline environment (BNG))*

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS 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*
MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*
RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*
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 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*
RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*
RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*
RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*
RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*
RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*
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 6073, *Segmented Pseudowire*
RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*
RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*
RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*
RFC 6718, *Pseudowire Redundancy*
RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*
RFC 6870, *Pseudowire Preferential Forwarding Status bit*
RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*

RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

Quality of Service (QoS)

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*

RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*

RFC 2598, *An Expedited Forwarding PHB*

RFC 3140, *Per Hop Behavior Identification Codes*

RFC 3260, *New Terminology and Clarifications for Diffserv*

Remote Authentication Dial In User Service (RADIUS)

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*

RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

RFC 2869, *RADIUS Extensions*

RFC 3162, *RADIUS and IPv6*

RFC 4818, *RADIUS Delegated-IPv6-Prefix Attribute*

RFC 5176, *Dynamic Authorization Extensions to RADIUS*

RFC 6911, *RADIUS attributes for IPv6 Access Networks*

RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*

Resource Reservation Protocol — Traffic Engineering (RSVP-TE)

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*

RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*

RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*

RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions (IF_ID RSVP_HOP object with unnumbered interfaces and RSVP-TE graceful restart helper procedures)*

RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*

RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*

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 Record Route Object (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 4950, *ICMP Extensions for Multiprotocol Label Switching*

RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 5712, *MPLS Traffic Engineering Soft Preemption*

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

Routing Information Protocol (RIP)

RFC 1058, *Routing Information Protocol*

RFC 2080, *RIPng for IPv6*

RFC 2082, *RIP-2 MD5 Authentication*

RFC 2453, *RIP Version 2*

Segment Routing (SR)

draft-filsfils-spring-segment-routing-policy-05, *Segment Routing Policy for Traffic Engineering*

draft-francois-rtgwg-segment-routing-ti-lfa-04, *Topology Independent Fast Reroute using Segment Routing*

draft-gredler-idr-bgp-ls-segment-routing-ext-03, *BGP Link-State extensions for Segment Routing*

draft-ietf-idr-segment-routing-te-policy-02, *Advertising Segment Routing Policies in BGP*

draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*

draft-ietf-mpls-spring-lsp-ping-02, Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing

draft-ietf-spring-conflict-resolution-05, Segment Routing MPLS Conflict Resolution

draft-ietf-spring-segment-routing-ldp-interop-09, Segment Routing interworking with LDP

Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)

ANSI T1.105.03, Jitter Network Interfaces

ANSI T1.105.06, Physical Layer Specifications

ANSI T1.105.09, Network Timing and Synchronization

ITU-T G.703, Physical/electrical characteristics of hierarchical digital interfaces

ITU-T G.707, Network node interface for the synchronous digital hierarchy (SDH)

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC)

ITU-T G.823, The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy

ITU-T G.824, The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy

ITU-T G.825, The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)

ITU-T G.841, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002

ITU-T G.957, Optical interfaces for equipments and systems relating to the synchronous digital hierarchy

Time Division Multiplexing (TDM)

ANSI T1.403, DS1 Metallic Interface Specification

ANSI T1.404, DS3 Metallic Interface Specification

Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions*, issued 09/2008

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*, issued 03/2003

ITU-T G.8261, *Timing and synchronization aspects in packet networks*, issued 04/2008

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC)*, issued 08/2007

ITU-T G.8264, *Distribution of timing information through packet networks*, issued 10/2008

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization*, issued 10/2010

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network*, issued 07/2014

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

Two-Way Active Measurement Protocol (TWAMP)

RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP) (server, unauthenticated mode)*

RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

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)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

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*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications* (estimating the interarrival jitter)

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

Wireless Local Area Network (WLAN) Gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses* (S2a roaming based on GPRS)

Customer Document and Product Support



Customer Documentation

[Customer Documentation Welcome Page](#)



Technical Support

[Product Support Portal](#)



Documentation Feedback

[Customer Documentation Feedback](#)

