



7210 Service Access System

Release 25.9.R1

7210 SAS-Mxp, R6, R12, S, Sx, T Quality of Service Guide

3HE 21177 AAAB TQZZA 01
Edition: 01
September 2025

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2025 Nokia.

Table of contents

List of tables.....	18
List of figures.....	22
1 Getting started.....	23
1.1 About this guide.....	23
1.1.1 Document structure and content.....	24
1.2 7210 SAS modes of operation.....	24
1.3 7210 SAS port modes.....	26
1.4 Nokia 7210 SAS-series services configuration process.....	29
1.5 Conventions.....	30
1.5.1 Precautionary and information messages.....	30
1.5.2 Options or substeps in procedures and sequential workflows.....	30
2 QoS policies.....	32
2.1 QoS policies overview.....	32
2.1.1 Overview of QoS policies on 7210 SAS-T in access-uplink mode.....	33
2.1.2 Overview of QoS policies on 7210 SAS-T in network mode.....	35
2.1.3 Overview of QoS policies on 7210 SAS-Mxp in network mode.....	38
2.1.4 Overview of QoS policies on 7210 SAS-R6 and 7210 SAS-R12.....	43
2.1.5 Overview of QoS policies on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE.....	50
2.2 Network and service QoS policies.....	55
2.2.1 Network QoS policies in network mode.....	55
2.2.1.1 Network QoS policy "ip-interface" type.....	56
2.2.1.2 Network QoS policy "port" type.....	58
2.2.2 Network QoS policies in access-uplink mode.....	59
2.2.3 Network queue QoS policies.....	61
2.2.3.1 Network queue policies in network mode.....	61
2.2.3.2 Network queue policies in access-uplink mode.....	65
2.2.4 Service ingress QoS policies.....	66
2.2.4.1 CAM-based classification.....	68
2.2.4.2 Table-based classification.....	71
2.2.4.3 Hierarchical ingress policing.....	71
2.2.5 Service egress QoS policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	72

2.3	Access ingress QoS policies.....	75
2.4	Access egress QoS policies.....	78
2.4.1	Access egress QoS policies on 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.....	78
2.4.2	Access egress QoS policies on 7210 SAS-Mxp.....	81
2.4.2.1	Access egress QoS policy for SAP-based queuing mode on 7210 SAS-Mxp.....	81
2.4.2.2	Access egress QoS policy for port-based queuing mode on 7210 SAS-Mxp.....	81
2.4.3	Access egress QoS policies on 7210 SAS-R6 and 7210 SAS-R12.....	82
2.4.3.1	Access egress QoS policies for SAP-based queuing mode.....	82
2.4.3.2	Access egress QoS policy for port-based queuing mode.....	83
2.4.4	Queue overrides for access egress QoS policies on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	85
2.4.4.1	Configuring access egress QoS policy queue override parameters.....	86
2.5	Remark policies on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (network mode).....	86
2.5.1	Egress port rate limiting.....	87
2.6	Forwarding classes.....	87
2.6.1	Forwarding class-to-queue ID map on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T.....	88
2.6.2	Forwarding class-to-queue ID map on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE.....	89
2.7	QoS policy entities.....	90
2.7.1	QoS policies for hybrid ports on 7210 SAS-T.....	91
2.7.2	QoS policies for hybrid ports on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE.....	92
2.7.3	QoS policies for hybrid ports on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	92
2.8	Meters/policers.....	93
2.8.1	Meter/policer parameters.....	93
2.8.1.1	Meter ID.....	93
2.8.1.2	Committed information rate for meters.....	93
2.8.1.3	Peak information rate for meters.....	94
2.8.1.4	Color-aware and color-blind policers.....	94
2.8.1.5	Adaptation rule for meters.....	95
2.8.1.6	Committed burst size (for meter/policers).....	97
2.8.1.7	Maximum burst size (for meter/policers).....	97
2.8.1.8	Meter counters.....	98
2.8.1.9	Meter modes.....	98
2.8.1.10	QoS overrides for meters/policers.....	98

2.9	Queue management.....	99
2.9.1	Queue parameters.....	99
2.9.1.1	Queue ID.....	100
2.9.1.2	CIR for queues.....	100
2.9.1.3	PIR for queues.....	100
2.9.1.4	Adaptation rule for queues.....	101
2.9.1.5	CBS and MBS for queues.....	104
2.9.2	Buffer pools.....	105
2.9.2.1	Buffer pools on 7210 SAS-T.....	106
2.9.2.2	Decommissioning ports with per port MBS pool.....	107
2.9.2.3	Buffer pools on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE.....	109
2.9.2.4	Buffer pools on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	110
2.9.3	Queue management policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12..	110
2.9.3.1	Queue management policy parameters.....	111
2.9.4	RED slopes in network and access-uplink mode.....	111
2.9.4.1	Tuning the shared buffer utilization calculation.....	113
2.9.4.2	Slope policies for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE devices.....	115
2.9.5	CPU queues.....	117
2.10	Schedulers.....	117
2.10.1	Scheduler modes on 7210 SAS-T.....	117
2.10.2	Port scheduler policies for 7210 SAS-T.....	119
2.10.3	Schedulers on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE.....	119
2.10.4	Schedulers on 7210 SAS-Mxp.....	119
2.10.5	Schedulers on 7210 SAS-R6 and 7210 SAS-R12.....	120
2.11	Configuration notes.....	120
3	Discard eligibility indicator-based (DEI-based) classification and marking.....	121
3.1	DEI-based classification.....	121
3.2	DEI-based marking.....	122
3.3	Configuration guidelines.....	123
4	Port level egress rate-limiting.....	124
4.1	Overview.....	124
4.1.1	Applications.....	124
4.1.2	Effect of port level rate-limiting on network queue functionality.....	124
4.2	Basic configurations.....	125

4.2.1	Modifying port level egress-rate command.....	125
4.2.2	Removing port level egress-rate command.....	126
4.2.2.1	Default egress-rate values.....	126
4.3	Port level egress-rate command reference.....	126
4.3.1	Command hierarchies.....	126
4.3.1.1	Configuration commands.....	126
4.3.1.2	Show commands.....	126
4.3.2	Command descriptions.....	127
4.3.2.1	Configuration commands.....	127
4.3.2.2	Show commands.....	128
5	SAP egress aggregate meter.....	135
5.1	Overview.....	135
5.1.1	Configuration notes.....	135
5.2	Basic configurations.....	135
5.3	SAP egress aggregate meter command reference.....	136
6	Frame-based accounting.....	137
6.1	Overview.....	137
6.1.1	Frame-based accounting.....	137
6.1.2	Effects of enabling ingress frame-based accounting on ingress meter functionality....	137
6.1.3	Effects of enabling egress frame-based accounting on network queue functionality....	137
6.1.4	Accounting and statistics.....	138
6.2	Basic configurations.....	138
6.2.1	Enabling and disabling frame-based accounting.....	138
6.2.1.1	Default frame-based accounting values.....	139
6.3	Frame-based accounting command reference.....	139
6.3.1	Command hierarchies.....	139
6.3.1.1	Configuration commands.....	139
6.3.2	Configuration commands.....	139
	frame-based-accounting.....	139
	egress-enable.....	140
	ingress-enable.....	140
7	Network QoS policies.....	142
7.1	Overview.....	142

7.1.1	Overview of network QoS policies in network mode.....	142
7.1.1.1	Overview of network QoS policies in access-uplink mode.....	142
7.2	Network QoS policy in network mode.....	143
7.2.1	Network QoS policy (ip-interface type) behaviour for MPLS LSPs.....	144
7.2.2	Basic configurations.....	144
7.2.3	Create a network QoS policy (ip-interface type) for network mode.....	144
7.2.3.1	Example for network QoS policy of ip-interface type.....	146
7.2.4	Configuring network QoS policy (port type) for network mode.....	146
7.2.5	Default network policy values available in network mode.....	149
7.2.6	Resource allocation for network QoS policy.....	153
7.2.6.1	Network QoS policies resource usage examples.....	154
7.3	Network QoS policy in access-uplink mode.....	162
7.3.1	Basic configurations.....	163
7.3.2	Configuring network policy for access-uplink mode.....	163
7.3.3	Default network policy values available in access-uplink mode.....	165
7.4	DSCP and dot1p marking for CPU self-generated traffic.....	166
7.4.1	QoS for self-generated (CPU) traffic on network interfaces for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	168
7.4.2	Default DSCP mapping table.....	169
7.5	Service management tasks.....	170
7.5.1	Deleting QoS policies.....	170
7.5.2	Remove a policy from the QoS configuration.....	170
7.5.3	Copying and overwriting network policies.....	170
7.5.4	Editing QoS policies.....	171
7.6	Network QoS policy command reference.....	171
7.6.1	Command hierarchies.....	171
7.6.1.1	Configuration commands for MPLS EXP profile map (7210 SAS platforms operating in network mode).....	171
7.6.1.2	Configuration commands (7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE configured in network mode).....	171
7.6.1.3	Configuration commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS- R12).....	172
7.6.1.4	Configuration commands (access-uplink mode).....	173
7.6.1.5	Self-generated traffic commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12).....	174
7.6.1.6	Operational commands (network mode or access-uplink mode).....	174
7.6.1.7	Show commands (network mode or access-uplink mode).....	174
7.6.1.8	Show commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12).....	174

7.6.2	Command descriptions.....	175
7.6.2.1	Configuration commands.....	175
7.6.2.2	Network QoS policy commands.....	176
7.6.2.3	Network QoS policy commands (7210 SAS-T in access-uplink mode).....	180
7.6.2.4	Network ingress QoS policy commands.....	182
7.6.2.5	Network egress QoS policy commands.....	197
7.6.2.6	Self-generated traffic commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12).....	207
7.6.2.7	Operational commands.....	209
7.6.2.8	Show commands.....	210
7.6.2.9	Show commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12).....	222
8	Network queue QoS policies.....	233
8.1	Overview.....	233
8.2	Basic configurations.....	233
8.2.1	Create a network queue QoS policy.....	233
8.2.2	Applying network queue policies.....	236
8.2.2.1	Applying network queue configuration in network mode.....	236
8.2.2.2	Applying network queue configuration in access-uplink mode.....	236
8.3	Default network queue policy values.....	237
8.4	Default network queue policy values for hybrid ports on 7210 SAS-R6 and 7210 SAS-R12...	242
8.5	Service management tasks.....	243
8.5.1	Deleting network queue QoS policies.....	243
8.5.2	Copying and overwriting network queue QoS policies.....	243
8.5.3	Editing network queue QoS policies.....	244
8.6	Network queue QoS policy command reference.....	244
8.6.1	Command hierarchies.....	244
8.6.1.1	Configuration commands for 7210 SAS-T (in network mode and access-uplink mode).....	245
8.6.1.2	Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	245
8.6.1.3	Configuration commands for 7210 SAS-Sx 1/10GE and 7210 SAS-Sx 10/100GE.....	245
8.6.1.4	Operational commands.....	246
8.6.1.5	Show commands.....	246
8.6.2	Command descriptions.....	246
8.6.2.1	Configuration commands.....	246

8.6.2.2	Network queue QoS policy commands.....	247
8.6.2.3	Network queue QoS policy queue commands.....	248
8.6.2.4	Operational commands.....	253
8.6.2.5	Show commands.....	254
9	Service ingress QoS policies.....	260
9.1	Overview.....	260
9.1.1	Default SAP ingress policy.....	260
9.1.2	SAP-ingress policy defaults.....	261
9.2	Resource allocation for SAP ingress policy.....	261
9.2.1	Use of index file by SAP QoS ingress policy.....	261
9.2.2	Use of the keyword "multipoint" for default meter "11".....	262
9.2.2.1	Example uses of the multipoint meter.....	262
9.2.3	Service ingress meter selection rules.....	264
9.2.3.1	Default policy.....	264
9.2.3.2	VPLS service without meter "11".....	264
9.2.3.3	VPLS service with meter "11".....	265
9.2.3.4	Epipe, IES, and VPRN services without PIM.....	265
9.2.3.5	IES and VPRN services with PIM/multicast and without meter "11".....	265
9.2.3.6	IES and VPRN services with PIM/multicast and meter "11".....	266
9.2.4	Service ingress policy configuration considerations.....	266
9.2.5	Resource allocation for service ingress QoS policies using CAM-based classification..	267
9.2.5.1	Resource configuration guidelines for service ingress QoS policies using CAM-based classification.....	270
9.2.6	Computation of resources used per SAP ingress policy for CAM-based classification..	270
9.2.6.1	Determining the number of classification entries.....	271
9.3	Table-based classification using dot1p and IP DSCP for assigning FC and profile on SAP ingress for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	274
9.3.1	IP DSCP and dot1p classification policy support.....	274
9.3.1.1	Default-FC assignment rules for SAPs in Layer 3 services.....	275
9.3.2	Precedence rules for DEI assignments on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	276
9.3.3	Creating an IP DSCP and dot1p classification policy.....	276
9.3.4	CAM resource usage for IP DSCP and dot1p classification policies on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	277
9.3.4.1	CAM resource allocation for table-based classification.....	278
9.3.5	Associating a DSCP or dot1p classification policy.....	278

9.3.5.1	Associating a classification policy with a SAP ingress QoS policy.....	278
9.3.5.2	Associating a classification policy with an Ethernet port.....	279
9.3.6	Assigning and enabling DSCP and dot1p classification policies to a SAP.....	280
9.3.6.1	Assigning and enabling policies to Epipe and VPLS SAPs.....	281
9.3.6.2	Assigning and enabling policies to IES and VPRN interface SAPs.....	282
9.3.6.3	Assigning policies to RVPLS SAPs.....	282
9.4	Service meter for SAP ingress (7210 SAS-Mxp).....	285
9.4.1	Default service meter policy.....	286
9.4.2	Resource usage for service meters.....	286
9.4.2.1	Examples for service meters with computation of resource usage.....	287
9.5	Calculating resources required for classification.....	290
9.5.1	Examples: calculating resources required for CAM-based classification.....	290
9.5.1.1	Example 1.....	291
9.5.1.2	Example 1a (default multipoint meter 11 is not used):.....	292
9.5.1.3	Example 2.....	294
9.5.1.4	Example 2a (default multipoint meter "11" is not used):.....	295
9.5.1.5	Example 3.....	297
9.5.1.6	Example 3a (default multipoint meter "11" is not used):.....	298
9.5.1.7	Example 4.....	300
9.5.1.8	Example 4a (default multipoint meter "11" is not used):.....	303
9.5.1.9	Example 5.....	305
9.5.1.10	Example 6.....	307
9.5.1.11	Example 7.....	309
9.5.1.12	Example 8.....	310
9.5.1.13	Example 9.....	311
9.5.1.14	Example 9a (default multipoint meter "11" is not used):.....	315
9.5.1.15	Example 10.....	320
9.5.1.16	Example 11.....	321
9.5.2	Examples: calculating resources required for IP DSCP table-based classification with CAM-based policing (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12).....	322
9.5.2.1	Example 1: Epipe, IES, and VPRN services using unicast traffic type.....	323
9.5.2.2	Example 2: VPLS using unicast and BUM meter with IES or VPRN using multicast.....	326
9.5.2.3	Example 3: VPLS service using unicast, broadcast, multicast, and unknown- unicast with additional FCs.....	328
9.5.2.4	Example 4: routed VPLS on access port using unicast, broadcast, multicast, and unknown-unicast with additional FCs.....	331

9.5.2.5	Example 5: routed VPLS service on a hybrid port using unicast, broadcast, multicast and unknown-unicast for some FCs.....	336
9.5.2.6	Example 6: routed VPLS on access port and hybrid port.....	340
9.6	Basic configurations.....	345
9.6.1	Create service ingress QoS policies.....	345
9.6.1.1	Service ingress QoS meter.....	346
9.6.1.2	Service ingress IP match criteria.....	347
9.6.1.3	Service ingress MAC match criteria.....	348
9.6.2	Applying service ingress policies.....	349
9.6.2.1	Epipe.....	349
9.6.2.2	VPLS.....	349
9.6.2.3	VPRN.....	350
9.6.2.4	IES.....	350
9.7	Service management tasks.....	351
9.7.1	Deleting QoS policies.....	351
9.7.1.1	Remove a QoS policy from service SAPs.....	351
9.7.2	Copying and overwriting QoS policies.....	351
9.7.3	Remove a policy from the QoS configuration.....	352
9.7.4	Editing QoS policies.....	353
9.8	Service ingress QoS policy command reference.....	353
9.8.1	Command hierarchy.....	353
9.8.1.1	Service ingress QoS policy commands.....	353
9.8.1.2	Table-based IP DSCP and dot1p classification policy commands for SAP ingress (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12).....	355
9.8.1.3	Service meter commands for SAP ingress (7210 SAS-Mxp only).....	355
9.8.1.4	Operational commands.....	356
9.8.1.5	Show commands.....	356
9.8.2	Command descriptions.....	356
9.8.2.1	Configuration commands.....	356
9.8.2.2	Operational commands.....	406
9.8.2.3	Show commands.....	408
10	Access egress QoS policies on 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.....	426
10.1	Overview.....	426
10.1.1	Basic configurations.....	426
10.1.1.1	Modifying access egress QoS queues.....	428

10.1.1.2	Applying access egress QoS policies.....	428
10.1.1.3	Default access egress QoS policy values.....	429
10.1.1.4	Deleting QoS policies.....	431
10.1.1.5	Removing a policy from the QoS configuration.....	431
10.2	Access egress QoS policy command reference.....	431
10.2.1	Command hierarchies.....	431
10.2.1.1	Configuration commands for 7210 SAS-T (access-uplink mode).....	431
10.2.1.2	Configuration commands for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE (network mode).....	432
10.2.1.3	Operational commands.....	432
10.2.1.4	Show commands.....	432
10.2.2	Configuration commands.....	433
10.2.2.1	Generic commands.....	433
10.2.2.2	Access egress queue QoS policy commands.....	443
10.2.2.3	Operational commands.....	447
10.2.2.4	Show commands.....	447
11	Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	452
11.1	Overview.....	452
11.1.1	Access egress QoS policy for SAP-based queuing mode.....	452
11.1.2	Access egress QoS policy for port-based queuing mode.....	453
11.1.3	Access egress QoS policy queue override.....	454
11.1.4	Basic configurations.....	455
11.1.4.1	Modifying access egress QoS queues on the 7210 SAS-R6 and 7210 SAS- R12.....	455
11.1.4.2	Applying access egress QoS policies on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	455
11.1.4.3	Editing QoS policies.....	456
11.1.4.4	Deleting QoS policies.....	456
11.1.4.5	Removing a policy from the QoS configuration.....	457
11.2	Access egress QoS policy command reference.....	457
11.2.1	Command hierarchies.....	457
11.2.1.1	Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS- R12.....	457
11.2.1.2	Show commands.....	457
11.2.2	Command descriptions.....	458
11.2.2.1	Generic commands.....	458

11.2.2.2	Access egress queue QoS policy commands.....	462
11.2.2.3	Show commands.....	468
12	Service egress policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	473
12.1	Overview.....	473
12.1.1	Basic configurations.....	473
12.1.2	Create a SAP egress policy.....	474
12.1.3	Editing QoS policies.....	476
12.2	Service egress policy command reference.....	476
12.2.1	Command hierarchies.....	476
12.2.1.1	Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	476
12.2.1.2	Operational commands.....	477
12.2.1.3	Show commands.....	477
12.2.2	Command descriptions.....	477
12.2.2.1	Configuration commands.....	477
12.2.2.2	Operational commands.....	487
12.2.2.3	Show Commands.....	488
13	QoS port scheduler policies for 7210 SAS-T.....	494
13.1	Overview.....	494
13.1.1	Configuring port scheduler policies.....	494
13.1.2	Basic configurations.....	494
13.1.2.1	Creating a QoS port scheduler policy.....	494
13.2	Service management tasks.....	495
13.2.1	Copying and overwriting scheduler policies.....	495
13.2.2	Editing QoS policies.....	496
14	Schedulers on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE.....	497
14.1	Configuring scheduler policies.....	497
14.2	QoS port scheduler policy command reference.....	498
14.2.1	Command hierarchies.....	498
14.2.1.1	Port scheduler policy configuration commands.....	498
14.2.1.2	Operational commands.....	498
14.2.1.3	Show commands.....	499
14.2.2	Command descriptions.....	499
14.2.2.1	Configuration commands.....	499

14.2.2.2	Port scheduler policy commands.....	500
14.2.2.3	Operational commands.....	502
14.2.2.4	Show commands.....	503
15	Schedulers on 7210 SAS-Mxp.....	508
15.1	Overview.....	508
15.2	Scheduling with SAP-based queues on access ports.....	508
15.3	Scheduling on network ports.....	510
15.4	Scheduling on hybrid port with SAP-based egress queues.....	510
15.4.1	Port-based scheduling and queuing on access ports.....	511
15.5	Scheduling on hybrid port with port-based SAP queues.....	512
16	Schedulers on 7210 SAS-R6 and 7210 SAS-R12.....	514
16.1	Scheduling with SAP-based queues on access ports.....	514
16.2	Scheduling on network ports.....	516
16.3	Scheduling on hybrid port with SAP-based egress queues.....	516
16.3.1	Port-based scheduling and queuing on access ports.....	517
16.4	Scheduling on hybrid port with port-based SAP queues.....	518
17	Slope QoS policies.....	519
17.1	Overview.....	519
17.1.1	Configuration guidelines.....	519
17.1.2	WRED support on 7210 SAS-T access-uplink and network mode, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.....	519
17.1.3	Basic configurations.....	519
17.1.3.1	Create a slope QoS policy.....	520
17.1.3.2	Applying slope policies.....	521
17.1.4	Default slope policy values.....	521
17.2	Service management tasks.....	523
17.2.1	Deleting QoS policies.....	523
17.2.1.1	Ports.....	523
17.2.1.2	Remove a policy from the QoS configuration.....	524
17.2.2	Copying and overwriting QoS policies.....	524
17.2.3	Editing QoS policies.....	525
17.3	Slope QoS policy command reference.....	526
17.3.1	Command hierarchies.....	526
17.3.1.1	Configuration commands.....	526

17.3.1.2	Operational commands.....	526
17.3.1.3	Show commands.....	526
17.3.2	Command descriptions.....	527
17.3.2.1	Configuration commands.....	527
17.3.2.2	Operational commands.....	534
17.3.2.3	Show commands.....	535
18	Queue management policies.....	539
18.1	Overview.....	539
18.1.1	Basic configurations.....	539
18.1.2	Service management tasks.....	540
18.1.2.1	Creating a queue management policy.....	540
18.1.2.2	Editing QoS policies.....	541
18.2	Queue management policy command reference.....	541
18.2.1	Command hierarchies.....	541
18.2.1.1	Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	541
18.2.1.2	Operational commands.....	542
18.2.1.3	Show commands.....	542
18.2.2	Command descriptions.....	542
18.2.2.1	Configuration commands.....	542
18.2.2.2	Operational commands.....	550
18.2.2.3	Show commands.....	550
19	Remark policies.....	554
19.1	Overview.....	554
19.1.1	Configuration guidelines.....	557
19.1.2	Basic configurations.....	558
19.1.2.1	Creating a remark policy.....	558
19.1.2.2	Editing QoS policies.....	560
19.2	Remark policy command reference.....	560
19.2.1	Command hierarchies.....	560
19.2.1.1	Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T in network mode.....	560
19.2.1.2	Operational commands.....	561
19.2.1.3	Show commands.....	561

19.2.2	Command descriptions.....	561
19.2.2.1	Configuration commands.....	561
19.2.2.2	Operational commands.....	573
19.2.2.3	Show commands.....	574
20	Access ingress QoS policies.....	583
20.1	Overview.....	583
20.1.1	Shared access ingress QoS policies on the 7210 SAS-Mxp.....	584
20.1.2	Meter allocation rules for shared access ingress QoS policies.....	584
20.1.3	Resource allocation for non-shared access ingress QoS policies.....	584
20.1.3.1	Using index file for access ingress QoS policies.....	585
20.1.3.2	Calculating the number of QoS resources for non-shared access ingress QoS policies.....	585
20.1.3.3	Calculating the number of meters or policers for non-shared access ingress QoS policies.....	586
20.1.3.4	Determining the number of resources allocated to non-shared access ingress QoS policies.....	586
20.1.3.5	Example of non-shared access ingress QoS policy resource calculations.....	586
20.1.4	Resource allocation for shared access ingress QoS policies on the 7210 SAS-Mxp.....	588
20.1.4.1	Calculating the number of QoS resources for shared access ingress QoS policies.....	589
20.1.4.2	Calculating the number of meters or policers for shared access ingress QoS policies.....	590
20.1.4.3	Determining the number of resources allocated to shared access ingress QoS policies.....	591
20.1.4.4	Examples of shared access ingress QoS policy resource calculations.....	591
20.1.5	Configuration guidelines for a port-based access ingress QoS policy.....	593
20.1.6	Basic configurations for non-shared access ingress QoS policies.....	594
20.1.6.1	Editing a non-shared access ingress QoS policy configuration.....	595
20.1.6.2	Removing a non-shared policy from the QoS configuration.....	595
20.1.6.3	Deleting a non-shared access ingress QoS policy.....	595
20.1.7	Basic configurations for shared access ingress QoS policies.....	596
20.1.7.1	Editing a shared access ingress QoS policy configuration.....	596
20.1.7.2	Removing a shared policy from the QoS configuration.....	597
20.1.7.3	Deleting a shared access ingress QoS policy.....	597
20.2	Access-ingress QoS policy command reference.....	597
20.2.1	Command hierarchies.....	597
20.2.1.1	Access-ingress QoS configuration commands.....	597

20.2.1.2	Show commands.....	599
20.2.2	Command descriptions.....	599
20.2.2.1	Generic commands.....	599
20.2.2.2	Access ingress QoS policy commands.....	601
20.2.2.3	Show commands.....	626
21	Standards and protocol support.....	634
21.1	BGP.....	634
21.2	Ethernet.....	636
21.3	EVPN.....	637
21.4	Fast Reroute.....	637
21.5	Internet Protocol (IP) — General.....	638
21.6	IP — Multicast.....	640
21.7	IP — Version 4.....	641
21.8	IP — Version 6.....	642
21.9	IPsec.....	643
21.10	IS-IS.....	644
21.11	Management.....	645
21.12	MPLS — General.....	648
21.13	MPLS — GMPLS.....	649
21.14	MPLS — LDP.....	649
21.15	MPLS — MPLS-TP.....	649
21.16	MPLS — OAM.....	650
21.17	MPLS — RSVP-TE.....	650
21.18	OSPF.....	651
21.19	Pseudowire.....	652
21.20	Quality of Service.....	653
21.21	RIP.....	653
21.22	Timing.....	653
21.23	VPLS.....	655

List of tables

Table 1: Supported modes of operation and configuration methods.....	25
Table 2: Supported port modes by mode of operation.....	27
Table 3: 7210 SAS platforms supporting port modes.....	28
Table 4: Configuration process.....	29
Table 5: QoS policy types and descriptions for 7210 SAS-T devices in access-uplink mode.....	34
Table 6: QoS policy types and descriptions for 7210 SAS-T devices in network mode.....	37
Table 7: QoS policy types and descriptions for 7210 SAS-Mxp devices in network mode.....	41
Table 8: QoS policy types and descriptions for 7210 SAS-R6 and 7210 SAS-R12.....	46
Table 9: QoS policy types and descriptions for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE devices in network mode.....	53
Table 10: Default network QoS policy (type = ip-interface) egress marking.....	57
Table 11: Default network QoS policy (type = ip-interface) EXP-to-FC mapping.....	57
Table 12: Default network QoS policy of type "port" egress marking.....	59
Table 13: Default network QoS policy of type "port": dot1p/DSCP-to-FC mapping.....	59
Table 14: Default network QoS policy used for egress marking on access-uplink ports.....	60
Table 15: Default network QoS policy used for dot1p to FC on access-uplink ports.....	61
Table 16: Default network queue policy definition (for 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T configured in network mode).....	62
Table 17: Default network queue policy definition (for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 configured in network mode).....	63
Table 18: Default network queue policy definition (for 7210 SAS-T configured in access-uplink mode).....	65
Table 19: Service ingress QoS policy match criteria for 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T in network mode.....	69
Table 20: Service ingress QoS policy criteria for 7210 SAS-T in access-uplink mode.....	70

Table 21: MAC match Ethernet frame types.....	70
Table 22: MAC match criteria frame type dependencies.....	70
Table 23: Default service ingress policy ID 1 definition.....	71
Table 24: Default service egress policy "default" definition.....	74
Table 25: Default access ingress policy ID 1 definition for 7210 SAS-Mxp (non-shared access ingress policy mode), 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.....	77
Table 26: Shared default access ingress policy ID 65536 definition for 7210 SAS-Mxp.....	77
Table 27: Default access egress policy ID 1 definition for 7210 SAS-T.....	79
Table 28: Default access egress QoS policy "default" definition for 7210 SAS-R6 and 7210 SAS-R12.....	84
Table 29: Default remarking policy for dot1p on 7210 SAS-R6 and 7210 SAS-R12.....	85
Table 30: Forwarding classes.....	87
Table 31: Forwarding class to queue-ID map.....	88
Table 32: Forwarding class-to-queue ID map for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE....	89
Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T.....	95
Table 34: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Sx 10/100GE.....	95
Table 35: Administrative rate example.....	96
Table 36: Supported hardware rates and CIR/PIR values for 7210 SAS-T and 7210 SAS-Sx/S 1/10GE devices.....	102
Table 37: Supported hardware rates and CIR/PIR values for 7210 SAS-Mxp.....	102
Table 38: Supported hardware rates for CIR and PIR values for 7210 SAS-R6 and 7210 SAS-R12.....	102
Table 39: Supported hardware rates and CIR/PIR values for 10-Gig port for all platforms.....	103
Table 40: Supported hardware rates and CIR/PIR values for 7210 SAS-Sx 10/100GE.....	103
Table 41: Default CBS and MBS values.....	105

Table 42: Default values for the default slope policy for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	111
Table 43: TAF impact on shared buffer average utilization calculation.....	114
Table 44: Default slope policy definition (for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE configured in network mode).....	116
Table 45: Minimum and maximum bandwidth meters example.....	118
Table 46: Output fields: specific port.....	129
Table 47: Network policy defaults for policy ip-interface type.....	149
Table 48: Default network QoS policy of ip-interface type, LSP EXP-to-FC mapping on Ingress.....	150
Table 49: Default CPU QoS values for DSCP and dot1p marking.....	167
Table 50: Default DSCP mapping table.....	169
Table 51: Default DSCP names to DSCP value mapping table.....	189
Table 52: Default class selector code points to DSCP value mapping table.....	189
Table 53: Output fields: QoS network policy.....	216
Table 54: Output fields: network QoS policyfor the 7210 SAS-R6 and 7210 SAS-R12.....	220
Table 55: Output fields: QoS MPLS LSP EXP profile map.....	222
Table 56: Output fields: QoS network DSCP.....	224
Table 57: Output fields: SGT-QoS application.....	228
Table 58: Output fields: SGT-QoS DSCP-to-FC mapping.....	232
Table 59: Output fields: network queue labels.....	256
Table 60: Output fields: network queue policy.....	258
Table 61: SAP-ingress policy defaults.....	261
Table 62: SAP ingress resource allocation and match criteria types.....	268
Table 63: Maximum CIR and PIR values for 7210 SAS platforms.....	393

Table 64: Output fields: dot1p classification policy.....	409
Table 65: Output fields: DSCP classification policy.....	412
Table 66: Output fields: SAP-ingress QoS policy.....	416
Table 67: Output fields: FC meter map policy.....	423
Table 68: Default FC marking values for 7210 SAS-T (access-uplink mode).....	430
Table 69: Output fields: QoS access egress.....	450
Table 70: Output fields: QoS access egress.....	471
Table 71: Output fields: SAP egress QoS policy (7210 SAS-Mxp).....	489
Table 72: Output fields: SAP egress QoS policy (7210 SAS-R6 and 7210 SAS-R12).....	491
Table 73: Output fields: port scheduler policy for 7210 SAS-T.....	504
Table 74: Output fields: port scheduler policy for 7210 SAS-Sx 1/10GE.....	505
Table 75: Output fields: association.....	507
Table 76: Output fields: QoS slope policy.....	537
Table 77: Output fields: queue management policy.....	552
Table 78: Summary of remark policy and attachment points for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE (network mode).....	555
Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	556
Table 80: Output fields: remark policy.....	581
Table 81: Output fields: access ingress QoS policy (shared and non-shared).....	631

List of figures

Figure 1: 7210 SAS traffic types operating in network mode.....

55

Figure 2: 7210 SAS-T traffic types for access-uplink mode.....

55

Figure 3: Traffic queuing model for forwarding classes.....

68

Figure 4: RED slope characteristics.....

113

Figure 5: Calculation for average shared buffer utilization.....

114

Figure 6: SAP egress scheduling.....

508

Figure 7: Scheduling on network ports.....

510

Figure 8: Hybrid port egress scheduling.....

511

Figure 9: Port-based scheduling and queuing.....

511

Figure 10: Scheduling hierarchy for a hybrid port with SAPs.....

513

Figure 11: SAP egress scheduling.....

514

Figure 12: Scheduling on network ports.....

516

Figure 13: Hybrid port egress scheduling.....

516

Figure 14: Port-based scheduling and queuing.....

517

Figure 15: Scheduling hierarchy for the hybrid port.....

518

1 Getting started

This chapter provides process flow information to configure Quality of Service (QoS) policies and provision services. It also provides an overview of the document organization and content, and describes the terminology used in this guide.

1.1 About this guide



Note:

Unless explicitly noted otherwise, this guide uses 7210 SAS-Dxp to refer to the 7210 SAS-Dxp 12p, 7210 SAS-Dxp 16p, and 7210 SAS-Dxp 24p platforms.

This guide describes the QoS functionality and provides information to configure QoS policies on the following 7210 SAS platforms, operating in one of the modes described in [Table 1: Supported modes of operation and configuration methods](#). If multiple modes of operation apply, they are explicitly noted in the topic.

- 7210 SAS-Mxp
- 7210 SAS-R6
- 7210 SAS-R12
- 7210 SAS-Sx/S 1/10GE
- 7210 SAS-Sx 10/100GE
- 7210 SAS-T

See [7210 SAS modes of operation](#) for information about the modes of operation supported by the 7210 SAS product family.



Note:

Unless explicitly noted otherwise, the phrase "Supported on all 7210 SAS platforms as described in this document" is used to indicate that the topic and CLI commands apply to all the 7210 SAS platforms in the following list, when operating in the specified modes only.

- network mode of operation
7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx 10/100GE, and 7210 SAS-T
- standalone mode of operation
7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE
- standalone-VC mode of operation
7210 SAS-Sx/S 1/10GE

If the topics and CLI commands are supported on the 7210 SAS-T operating in the access-uplink mode, it is explicitly indicated, where applicable.

1.1.1 Document structure and content

This guide uses the following structure to describe routing protocols and route policies content.



Note:

This guide generically covers Release 25.x.Rx content and may include some content that will be released in later maintenance loads. See the *7210 SAS Software Release Notes 25.x.Rx*, part number 3HE 21188 000x TQZZA, for information about features supported in each load of the Release 25.x.Rx software.

- This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow. Each chapter describes a software area and provides CLI syntax and command usage to configure parameters for the functional area.
- Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.
- Unless explicitly noted, the CLI commands and their configuration is similar for both [network](#) and [access-uplink](#) operating modes for features applicable to both modes of operation.

1.2 7210 SAS modes of operation

Unless explicitly noted, the phrase "mode of operation" and "operating mode" refers to the current operating mode of the 7210 SAS router. Each operating mode provides configuration access to a specific group of CLI commands.



Note:

Not all CLI commands are supported on all 7210 SAS platforms in all modes of operation. Users can only configure CLI commands supported by the current operating mode of the router. See the *7210 SAS Software Release Notes 25.x.Rx*, part number 3HE 21188 000x TQZZA, and to the appropriate 7210 SAS software user guide for information about features and capabilities supported by a 7210 SAS platform when operating in a specific mode.

The following modes of operation are supported by the 7210 SAS product family.

- **access-uplink**

In the access-uplink operating mode, the 7210 SAS router uplinks to the network using Layer 2 Ethernet VLAN switching (without IP/MPLS).

Platforms Supported: 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C, and 7210 SAS-T

- **network**

In the network operating mode, the 7210 SAS router uses IP/MPLS uplinks to the network. The IP routing protocols and MPLS functionality is available; see the appropriate 7210 SAS software user guide for more information about supported features.

Platforms Supported: 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C, 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T

- **satellite**

In the satellite operating mode, the 7210 SAS platform uses high-capacity uplinks (for example, 10GE ports on the 7210 SAS-Mxp and 100GE ports on the 7210 SAS-Sx 10/100GE) to connect to the 7750 SR host. The 7210 SAS router is managed by the 7750 SR host. There is no direct CLI access to the satellite node, and all services and protocols are configured on the host.

See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for boot options to configure the [satellite](#) mode of operation on the router. See the 7750 SR software user guides for information about service and protocol provisioning, and operating the 7210 SAS router in [satellite](#) mode.

Platforms Supported: 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE

- **standalone**

In the standalone operating mode, the 7210 SAS platform supports IP/MPLS uplinks. It is operated and managed independently.

The functionality and features available on the standalone 7210 SAS platform are similar to the [network](#) operating mode. The standalone mode is primarily used to differentiate between a node being managed by the 7750 SR host (in the [satellite](#) operating mode), and a node managed independently (standalone operating mode).

Platforms Supported: 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE

- **standalone-VC**

In the standalone-VC operating mode, a set of 7210 SAS devices are stacked to provide larger 1GE/10GE port density and control-plane redundancy. The stack of nodes is provisioned and managed as a single chassis, and not as individual nodes.

The functionality and features available on the 7210 SAS platform are similar to the [network](#) operating mode, with additional capabilities, such as control-plane redundancy with non-stop routing and non-stop services.

Platforms Supported: 7210 SAS-Sx/S 1/10GE

For 7210 SAS platforms that support multiple explicit modes of operation ([Table 1: Supported modes of operation and configuration methods](#)), the operating mode must be configured in the Boot Option File (BOF) to ensure the router boots up in the specified mode. For example, the 7210 SAS-T supports access-uplink and network modes of operation, and the 7210 SAS-Sx/S 1/10GE supports satellite, standalone, and standalone-VC mode of operations. In some cases, the 7210 SAS router operates in a specific mode implicitly, and explicit configuration is not required.

See the appropriate *Basic System Configuration Guide* for boot options and information about how to boot the 7210 SAS platform in a specific operating mode.

The following table lists the supported modes of operation and the configuration methods for the 7210 SAS platforms. Unless explicitly noted otherwise, the operating mode is supported on all variants of the specific 7210 SAS platform.

Table 1: Supported modes of operation and configuration methods

7210 SAS platform	Mode of operation and configuration method				
	Network	Access-uplink	Standalone	Standalone-VC	Satellite
7210 SAS-D		Implicit	Implicit		
7210 SAS-Dxp		Implicit	Implicit		

7210 SAS platform	Mode of operation and configuration method				
	Network	Access-uplink	Standalone	Standalone-VC	Satellite
7210 SAS-K 2F1C2T		Implicit	Implicit		
7210 SAS-K 2F6C4T ¹	Port Mode Configuration ²	Port Mode Configuration ²	Implicit		
7210 SAS-K 3SFP+ 8C ¹	Port Mode Configuration ²	Port Mode Configuration ²	Implicit		
7210 SAS-Mxp	Implicit ³		Explicit BOF Configuration		Explicit BOF Configuration
7210 SAS-R6 ⁴	Implicit		Implicit		
7210 SAS-R12 ⁴	Implicit		Implicit		
7210 SAS-Sx/S 1/10GE	Implicit ³		Explicit BOF Configuration	Explicit BOF Configuration	Explicit BOF Configuration
7210 SAS-Sx 10/100GE	Implicit ³		Explicit BOF Configuration		Explicit BOF Configuration
7210 SAS-T	Explicit BOF Configuration	Explicit BOF Configuration	Implicit		

1.3 7210 SAS port modes

Unless explicitly noted, the phrase “port mode” refers to the current port configuration of the 7210 SAS node. The 7210 SAS platform supports the configuration of the following port modes.

- **access port mode**

Access ports are configured for customer-facing traffic if Service Access Points (SAPs) are required. The appropriate encapsulation type must be configured to distinguish the services on the port; services are configured on the port based on the encapsulation value.

Access ports can be configured on all the 7210 SAS platforms.

- **access-uplink port mode**

- ¹ By default, the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C boot up in the [network](#) mode of operation. These platforms also allow the use of [access-uplink port mode](#) (without explicit BOF configuration), which provides the option to use Layer 2 uplinks instead of IP/MPLS uplinks to the network core, similar to the 7210 SAS-K 2F1C2T router.
- ² See section [7210 SAS port modes](#) for information about port mode configuration
- ³ Implicitly operates in [network](#) mode when [standalone](#) mode of operation is configured
- ⁴ Supports MPLS uplinks only and implicitly operates in network mode

Access-uplink ports provide native Ethernet connectivity in service provider transport or in an infrastructure network. With this option, the encap-type can be configured to only QinQ. Access-uplink SAPs, which are QinQ SAPs, can only be configured on an access-uplink port to allow the operator to differentiate multiple services being carried over a single uplink port.

This is the default port mode of a 7210 SAS node in the [access-uplink](#) mode of operation.

- **network port mode**

Network ports are configured for network-facing traffic in the service provider transport or infrastructure network, and provide IP/MPLS uplinks.

This is the default port mode of a 7210 SAS node in the [network](#) or [standalone](#) mode of operation.

- **hybrid port mode**

Hybrid ports are configured for access and network facing traffic, and allow a single port to operate in both access and network modes.

Port modes available for configuration on a 7210 SAS node are determined by the current mode of operation of the router.



Note:

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C are unique; all port modes listed in [Table 2: Supported port modes by mode of operation](#) are available for configuration on the router, regardless of the current mode of operation.

The following table lists the port mode configuration support per 7210 SAS mode of operation.

Table 2: Supported port modes by mode of operation

Mode of operation	Supported port mode			
	Access	Network	Hybrid	Access-uplink
Access-uplink	✓			✓
Network	✓	✓	✓	
Satellite ⁵				
Standalone	✓	✓	✓	
Standalone-VC	✓	✓	✓	

The following table lists the port mode configuration supported by the 7210 SAS product family. See the appropriate *Interface Configuration Guide* for more information about configuring the port modes for a specific platform.

⁵ Port modes are configured on the 7750 SR host and managed by the host.

Table 3: 7210 SAS platforms supporting port modes

Platform	Port mode			
	Access	Network	Hybrid	Access-uplink
7210 SAS-D	Yes	No	No	Yes
7210 SAS-Dxp	Yes	No	No	Yes
7210 SAS-K 2F1C2T	Yes	No	No	Yes
7210 SAS-K 2F6C4T	Yes	Yes	Yes	Yes
7210 SAS-K 3SFP+ 8C	Yes	Yes	Yes	Yes
7210 SAS-Mxp	Yes	Yes	Yes	No
7210 SAS-R6 IMM-b (IMMv2)	Yes	Yes	Yes	No
7210 SAS-R6 IMM-c 100GE (IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-R12 IMM-b	Yes	Yes	Yes	No
7210 SAS-R12 IMM-c 100GE (IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-Sx/S 1/10GE	Yes	Yes	Yes	No
7210 SAS-Sx 10/100GE	Yes	Yes	Yes	No
7210 SAS-T	Yes	Yes ⁶	Yes ⁷	Yes ⁸

⁶ Network ports are supported only if the node is operating in network mode.

⁷ Hybrid ports are supported only if the node is operating in network mode.

⁸ Access-uplink ports are supported only if the node is operating in access-uplink mode.

1.4 Nokia 7210 SAS-series services configuration process

The following table 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 4: Configuration process

Area	Task	Chapter
Policy configuration	Configuring QoS Policies	
	• DEI classification and marking	Discard eligibility indicator-based (DEI-based) classification and marking
	• egress Rate	Port level egress rate-limiting
	• accounting mode	Frame-based accounting
	• network	Network QoS policies
	• network queue	Network queue QoS policies
	• SAP ingress	Service ingress QoS policies
	• access egress	Access egress QoS policies on 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE
	• access egress for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12	Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12
	• service egress	Service egress policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12
	• port scheduler	QoS port scheduler policies for 7210 SAS-T
	• scheduler policies	Schedulers on 7210 SAS-Mxp
	• slope	Slope QoS policies
	• queue management policies	Queue management policies
	• remark policies	Remark policies
	• access ingress	Access ingress QoS policies
Reference	• list of IEEE, IETF, and other proprietary entities	Standards and protocol support

1.5 Conventions

This section describes the general conventions used in this guide.

1.5.1 Precautionary and information messages

The following information symbols are used in the documentation.



DANGER: Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.



WARNING: Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.



Caution: Caution indicates that the described activity or situation may reduce your component or system performance.



Note: Note provides additional operational information.



Tip: Tip provides suggestions for use or best practices.

1.5.2 Options or substeps in procedures and sequential workflows

Options in a procedure or a sequential workflow are indicated by a bulleted list. In the following example, at step 1, the user must perform the described action. At step 2, the user must perform one of the listed options to complete the step.

Example: Options in a procedure

1. User must perform this step.
2. This step offers three options. User must perform one option to complete this step:
 - This is one option.
 - This is another option.
 - This is yet another option.

Substeps in a procedure or a sequential workflow are indicated by letters. In the following example, at step 1, the user must perform the described action. At step 2, the user must perform two substeps (a. and b.) to complete the step.

Example: Substeps in a procedure

1. User must perform this step.
2. User must perform all substeps to complete this action:
 - a. This is one substep.

- b.** This is another substep.

2 QoS policies

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



Note:

The terms "meter" and "policer" are used interchangeably in this guide.

2.1 QoS policies overview

The 7210 SAS devices are designed with ingress and egress QoS mechanisms to support multiple services for each physical port. The 7210 SAS devices provide extensive and flexible capabilities to classify, police, queue, shape, and mark traffic.



Note:

Not all QoS capabilities are supported on all 7210 SAS platforms. The following chapters describe what is supported on different 7210 SAS platforms.

In the Nokia service router service model, a service is provisioned on the provider-edge (PE) equipment. Service data is encapsulated and then sent in a service tunnel to the far-end 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 behave like a Layer 2 path to the service data; however, the data is really traversing an IP or IP/MPLS core. The tunnel from one edge device to the other edge device is provisioned with encapsulation, and the services are mapped to the tunnel that most appropriately supports the service needs.

The 7210 SAS supports the following FCs, internally named: Network-Control, High-1, Expedited, High-2, Low-1, Assured, Low-2, and Best-Effort. See [Forwarding classes](#) for more information about the FCs.

The 7210 SAS supports the use of different types of QoS policies to handle the specific QoS needs at each point in the service delivery model within the device. QoS policies are defined in a global context in the 7210 SAS and only take effect when the policy is applied to an entity.

QoS policies are uniquely identified with a policy ID number or name. Typically, Policy ID 1 or Policy ID "default" is reserved for the default policy, which is used if no policy is explicitly applied. There are a few instances where the default QoS policy uses a different ID.

The QoS policies supported on the 7210 SAS can be divided into the following main types:

- QoS policies that are used for classification, defining metering and queuing attributes, and defining marking behavior.
- Slope policies that are used to define default buffer allocations and weighted random early detection (WRED) slope definitions.
- Port scheduler policies, SAP ingress and egress policies, and network and network-queue policies that determine how queues are scheduled.

2.1.1 Overview of QoS policies on 7210 SAS-T in access-uplink mode

When configured to operate in access-uplink mode, 7210 SAS-T QoS policies are applied on service ingress, access port egress, and access-uplink port ingress and egress. These policies allow users to configure the following:

- classification rules for how traffic is mapped to FCs
- FC association with meters and meter parameters used for policing (rate limiting)
- queuing parameters for shaping and buffer allocation
- QoS marking and interpretation

Several types of QoS policies exist:

- **service ingress policies for access SAP ingress**

Service ingress QoS policies are applied to the customer-facing SAPs. Traffic that enters through the SAP is classified to map it to an FC. FCs are associated with a meter or policer on ingress. The mapping of traffic to meters can be based on combinations of customer QoS marking (IEEE 802.1p bits), IP criteria, and MAC criteria. The characteristics of the FC meters are defined within the policy with regard to the number of FC meters for unicast traffic and the meter characteristics (such as CIR, PIR, and so on). Each of the FCs can be associated with different unicast parameters.

A service ingress QoS policy also defines up to three (3) meters per FC to be used for multipoint traffic for multipoint services. There can be up to 32 meters in total per service ingress QoS policy. In the case of the VPLS, four forwarding types (which is not to be confused with forwarding classes) are supported: 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.

- **access egress policies for access port egress**

An access egress policy is analogous to a SAP-egress policy, as defined in the 7x50 SR series of products. The difference is the point of attachment. An access egress policy is applied on the physical port as opposed to the logical port (SAP) for SAP-egress policy. It applies to all the SAPs on the port. An access egress QoS policy maps the traffic egressing the customer facing ports into various queues and marks the traffic accordingly. The FCs are mapped to the queues. There are 8 (eight) queues at the port level. FC-to-queue mapping is not configurable. The number of queues is not user-configurable and software always allocates 8 (eight) queues at the port level. An access egress policy also defines how to remark the FC-to-packet header bits (for example, IEEE 802.1p bits in the Layer 2 VLAN header, and so on).

- **network policies for access-uplink port, ingress and egress**

Network queue policies are applied on egress of access-uplink ports when operating in access-uplink mode. The policies define the FC queue characteristics for these entities. The FCs are mapped to the queues. There are 8 (eight) queues at the port level. FC-to-queue mapping is system-defined and not user-configurable. The number of queues is not user-configurable and software always allocates 8 (eight) queues at the port level.

For devices configured to operate in access-uplink mode, network QoS policies apply to access-uplink ports. Access-uplink ports in access-uplink mode are analogous to network ports in network mode. On ingress, the policy applied to an access-uplink port maps incoming dot1p values to FC and profile state for the traffic received from the core network. On egress, the policy maps FC and profile state to packet header values (for example, IEEE 802.1p value in the Layer 2 header) for traffic to be transmitted into the core network.

- **slope policies**

Slope policies are applied to the egress queues on the access, network, and hybrid ports. These policies define the WRED congestion management attributes, such as drop probability and thresholds for high-profile and low-profile traffic.

- **network-queue policies for access-uplink port, egress**

- **port-scheduler policies for access port and access-uplink port egress**

Service ingress, access egress, and network QoS policies are defined with a scope of either **template** or **exclusive**. Template policies can be applied to multiple entities (such as SAPs and ports); exclusive policies can only be applied to a single entity.

One service ingress QoS policy can be applied to a specific SAP. An access egress policy can be applied to an access port. One access egress QoS policy can be applied to the access port. One network QoS policy can be applied to an access-uplink port when operating in access-uplink mode. A network QoS policy defines both ingress and egress behavior. One network queue policy can be applied to an access-uplink port. If no QoS policy is explicitly applied to a SAP, port, or interface, a default QoS policy is applied.

The following table describes the major functions performed by the QoS policies.



Note:

Not all policies are supported on all platforms. See the following sections and chapters for more information.

Table 5: QoS policy types and descriptions for 7210 SAS-T devices in access-uplink mode

Policy type	Applied at...	Description	Section
Service Ingress	SAP ingress	<ul style="list-style-type: none"> • defines up to 16 FC meters and meter parameters for traffic classification • defines match criteria to map flows to the meters based on any one of the criteria 	Service ingress QoS policies
Access Egress	Access port	<ul style="list-style-type: none"> • defines up to 8 FC queues and queue parameters for traffic classification • maps FCs to the queues • defines FC to remarking values (for example: dot1p, and so on) • defines CIR levels and PIR weights that determine how the queue gets prioritized by the scheduler 	Access egress QoS policies on 7210 SAS-Mxp
Network	Access-uplink port	<ul style="list-style-type: none"> • used for classification/marketing of IP packets • at ingress, defines dot1p to FC mapping and 8 meters • at egress, defines FC to remarking values (for example, dot1p) 	Network QoS policies in network mode
Network Queue	Access-uplink port	<ul style="list-style-type: none"> • defines FC mappings to network queues and queue characteristics for the queues 	Network queue policies in network mode

Policy type	Applied at...	Description	Section
Slope	Access ports and access-uplink ports	<ul style="list-style-type: none"> enables or disables the high-slope, low-slope, and non-TCP parameters within the egress pool 	Slope policy parameters
Port scheduler	Access ports and access-uplink ports	<ul style="list-style-type: none"> defines the parameters for the port scheduler 	Port scheduler policies for 7210 SAS-T

2.1.2 Overview of QoS policies on 7210 SAS-T in network mode

QoS policies are applied on service ingress, access port egress, network port ingress and egress, and on network IP interface ingress when configured to operate in network mode.

These policies allow user to configure the following:

- classification rules to map traffic to FCs
- FC association with meters and meter parameters used for policing (rate-limiting)
- queuing parameters for shaping
- QoS marking and interpretation

Several types of QoS policies exist:

- **service ingress policies for access SAP ingress**

Service ingress QoS policies are applied to customer-facing SAPs. Traffic that enters through the SAP is classified to map it to an FC. FCs are associated with meters or policers on ingress. The mapping of traffic to meters can be based on combinations of customer QoS marking (IEEE 802.1p bits), IP criteria, and MAC criteria. The characteristics of the FC meters are defined in the policy as to the number of FC meters for unicast traffic and the meter characteristics (such as CIR, PIR, and so on). Each of the FCs can be associated with different unicast parameters. A service ingress QoS policy also defines up to three (3) meters per FC to be used for multipoint traffic for multipoint services. There can be up to 32 meters in total per service ingress QoS policy. In the case of VPLS, four forwarding types (which are not to be confused with FCs) are supported: 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.

- **access egress policies for access port egress**

An access egress policy is analogous to a SAP-egress policy, as defined in the 7x50 SR series of products. The difference is the point of attachment. An access egress policy is applied on the physical port as opposed to the logical port (SAP) for SAP egress policies. It applies to all the SAPs on the port. An access egress QoS policy maps the traffic egressing on the customer-facing ports into various queues and marks the traffic. The FCs are mapped to the queues. There are 8 (eight) queues at the port level. FC-to-queue mapping is static and not configurable. The number of queues is not user-configurable and the software allocates 8 (eight) queues at the port level. An access egress policy also defines how to remark the FC-to-packet header bits (for example, IEEE 802.1p bits in the Layer 2 VLAN header).

- **network policies for network and hybrid port, ingress and egress**

For devices configured to operate in network mode, there are two types of network QoS policies, one applied to a network IP interface and the other is applied to a network port. Network QoS policies are

applied to IP interfaces. On ingress, the policy applied to an IP interface maps incoming MPLS LSP EXP values to FC and profile state for the traffic received from the core network. On egress, the policy maps FC and profile state to MPLS LSP EXP values for traffic to be transmitted into the core network. The network policy applied to a network port maps incoming IP packets, DSCP or dot1p values, to the FC and the profile state for the traffic received from the core network. On egress, the policy maps FC and profile state to DSCP and/or dot1p values for IP traffic to be transmitted into the core network.

- **network-queue policies for network and hybrid port egress**

Network queue policies are applied on egress to network ports when operating in network mode. The policies define the FC queue characteristics for these entities. The FCs are mapped to the queues. There are 8 (eight) queues at the port level. FC-to-queue mapping is static and not configurable. The number of queues is not user-configurable and the software allocates 8 (eight) queues at the port level.

- **port-scheduler policies for access port, network port and hybrid port egress**

Port scheduler policies are applied on egress for access, network, and hybrid ports. These policies allow the user to define queue scheduling attributes, such as strict-priority queuing and weighted queuing.

- **slope policies**

Slope policies are applied to the egress queues on the access, network, and hybrid ports. These policies define the WRED congestion management attributes, such as drop probability and thresholds for high-profile and low-profile traffic.

- **remark policies for access port, network port, and hybrid port egress marking**

Remark policies are applied to access ports/SAPs, network ports/IP interfaces, and hybrid ports/SAPs/IP interfaces to configure the marking values for different forwarding classes and profiles. These policies enable marking of packet header QoS fields for packets that are forwarded out of the port, SAP, or IP interface. See [Remark policies](#) for more information.

Service ingress, access egress, and network QoS policies are defined with a scope of either **template** or **exclusive**. Template policies can be applied to multiple entities (such as SAPs and ports) whereas exclusive policies can only be applied to a single entity.

One service ingress QoS policy can be applied to a specific SAP access egress policy. One access egress QoS policy can be applied to the access port. One network QoS policy can be applied to a specific IP interface or network port, based on the type of network QoS policy, when operating in network mode. One network QoS policy can be applied to an access-uplink port when operating in access-uplink mode. A network QoS policy defines both ingress and egress behavior. One network queue policy can be applied to the network port or a access-uplink port.

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

The following table describes the major functions performed by QoS policies.



Note:

Not all policies are supported on all platforms. See the following sections and chapters for more information.

Table 6: QoS policy types and descriptions for 7210 SAS-T devices in network mode

Policy type	Applied at...	Description	Section
Service Ingress	SAP ingress	<ul style="list-style-type: none"> defines up to 16 FC meters and meter parameters for traffic classification defines match criteria to map flows to the meters based on any one of the criteria 	Service ingress QoS policies
Access Egress	Access port	<ul style="list-style-type: none"> defines up to 8 FC queues and queue parameters for traffic classification maps FCs to the queues defines FC-to-remarking values defines CIR levels and PIR weights, which determine how the scheduler prioritizes the queue 	Service egress QoS policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12
Network (of type "ip-interface")	IP interface	<ul style="list-style-type: none"> used for classification or marking of MPLS packets at ingress, defines MPLS LSP-EXP-to-FC mapping and 12 meters used by FCs at egress, defines FC-to-MPLS LSP-EXP marking 	Network QoS policies in network mode
Network (of type "port")	Network and hybrid ports	<ul style="list-style-type: none"> used for classification or marking of IP packets at ingress, defines DSCP or dot1p to FC mapping and 8 meters at egress, defines FC-to-DSCP or dot1p marking or both 	Network QoS policies in network mode
Network Queue	Network ports and hybrid ports	<ul style="list-style-type: none"> defines FC mappings to network queues, and queue characteristics for the queues 	Network queue policies in network mode
Slope	Access ports, network ports and hybrid ports	<ul style="list-style-type: none"> enables or disables the high-slope, low-slope, and non-TCP parameters within the egress pool 	Slope policy parameters
Port scheduler	Access ports, Network ports and Hybrid ports	<ul style="list-style-type: none"> defines the parameters for the port scheduler 	Port scheduler policies for 7210 SAS-T
Remark (Only on 7210 SAS-T)	Network ports, access	<ul style="list-style-type: none"> applied at egress, defines the FC-to-priority bits (DSCP or dot1p or EXP) marking 	Remark policies on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-

Policy type	Applied at...	Description	Section
	ports, and hybrid ports		R12, 7210 SAS-S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (network mode)

2.1.3 Overview of QoS policies on 7210 SAS-Mxp in network mode

QoS policies are applied on service ingress, access port egress, network port ingress and egress, and network IP interfaces ingress when configured to operate in network mode.

These policies allow users to configure the following:

- classification rules for how traffic is mapped to FCs
- FC association with meters and meter parameters used for policing (rate-limiting)
- queuing parameters for shaping
- QoS marking/interpretation

Several types of QoS policies exist:

- **service ingress policies for access SAP ingress**

Service ingress QoS policies are applied to customer-facing SAPs. Traffic that enters through the SAP is classified to map to an FC. FCs are associated with meters or policers on ingress. Mapping traffic meters can be based on combinations of customer QoS marking (IEEE 802.1p bits), IP criteria, and MAC criteria. The policy defines the number of FC meters for unicast traffic and other meter characteristics (such as CIR, PIR, and so on). Each FC can be associated with different unicast parameters.

A service ingress QoS policy also defines up to three meters per FC for multipoint traffic for use with multipoint services. The system supports the configuration of up to 32 meters per service ingress QoS policy. In the case of VPLS, four forwarding types (not to be confused with forward classes) are supported: unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are flooded to all destinations within the service, while the unicast forwarding type is handled in a point-to-point fashion within the service.

- **service egress policies for access SAP egress**

Service egress QoS policies are applied to SAPs and map FCs to service egress queues for a service. The system allocates a maximum of eight queues per SAP for the eight FCs; this is not user configurable. All traffic types (unicast and BUM) share the same service egress queue. A service egress QoS policy defines the FC queue characteristics and how the FC to priority bits in the packet header are remarked (for example, IEEE 802.1p bits in the Ethernet VLAN tag) in the customer traffic.

On the 7210 SAS-Mxp, the user has a per-node or per-chassis option of configuring the SAP-based or access port-based egress queuing mode. The SAP-based egress queuing mode (**port-scheduler-mode** command disabled) uses a service egress QoS policy with the capability to use eight egress queues per SAP. The access port-based egress queuing mode (**port-scheduler-mode** command enabled) uses an access egress policy with the capability to use eight egress queues for all SAPs configured on the access port.

A service egress QoS policy or access egress policy also defines how to remark the FC-to-IEEE 802.1p bits in the customer traffic.

- **access egress policies for access port egress**

An access egress policy is applied to all SAPs on the physical port (and not the logical port (SAP)) for SAP-egress policies. It applies to all the SAPs on the port. Access egress policies provide different capabilities based on the egress queuing mode applied to the node. In SAP-based egress queuing mode, an access egress policy defines the remarking of the FC-to-packet header bits. For example, IEEE 802.1p bits in the Layer 2 VLAN header.

In port-based egress queuing mode, the access egress policy, in addition to remarking, is used to define the queuing and scheduling behavior for the port-based egress queues. Access egress QoS policies are applied to ports and map forwarding classes (FCs) to port egress queues. The system allocates a maximum of eight queues per port for the eight FCs. The allocation is not user-configurable. All traffic types, (unicast and BUM traffic types) share the same queue on port egress.

An access egress QoS policy defines the FC queue characteristics and the remarking of the FC to priority bits in the packet header (for example, IEEE 802.1p bits in the Ethernet VLAN tag) in the customer traffic.

- **access ingress policies for access port ingress**

An access ingress policy is applied to the physical port instead of the SAP; the policy applies to all SAPs configured on the specific access port. At ingress, the access ingress QoS policy uses dot1p, DEI with dot1p, or IP DSCP values to assign an FC and profile to traffic, which facilitates the classification of traffic received on the access port.

An option is provided to share an access ingress policy across multiple access ports so that the classification policy and the policer rate applies to all access SAPs configured across all access ports that share the policy. In shared mode, the access ingress policy provides an option to use dot1p, DEI, IP DSCP, and IP criteria (with or without a port range) classification entries to assign an FC and profile to traffic received on the access port.

The FC assigned using the classification entry is associated with meters at ingress and allows the user to define up to one meter per FC for unicast traffic, and up to one meter per FC for multipoint traffic (broadcast, multicast, and unknown-unicast) for multipoint services. The system supports up to 16 meters per access ingress QoS policy.



Note:

This policy is available only when the node is operating in **sap-scale mode high**. See the *7210 SAS-Mxp, S, Sx, T Services Guide* and *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information about the **sap-scale-mode** command.

- **network policies for network and hybrid port, ingress and egress**

For devices configured to operate in network mode, two types of network QoS policies are supported: one that is applied to a network IP interface and the other to a network port. Network QoS policies are applied to IP interfaces. On ingress, the policy applied to an IP interface maps incoming MPLS LSP EXP values to FC and profile state for traffic received from the core network. On egress, the network policy maps FC and profile state to MPLS LSP EXP values for traffic transmitted into the core network. The network policy applied to a network port maps incoming IP packets, DSCP, or dot1p values to the FC and the profile state for the traffic received from the core network. On egress, the network policy maps FC and profile state to DSCP or dot1p values for IP traffic transmitted into the core network.

- **network queue policies for network and hybrid port, egress**

Network queue policies are applied on egress to network ports when operating in network mode. The policies define the FC queue characteristics for these entities. The FCs are mapped to the queues. The FC-to-queue mapping is static and not configurable. The number of queues is not user-configurable and the software allocates eight queues at the port level.

- **port scheduler policies for access port, network port and hybrid port egress**

Port scheduler policies are applied on egress for access, network, and hybrid ports. These policies allow the user to define queue scheduling attributes, such as strict-priority queuing and weighted queuing.

- **slope policies**

Slope policies are applied to the egress queues on the access, network, and hybrid ports. These policies define the WRED congestion management attributes, such as drop probability and thresholds for high-profile and low-profile traffic.

- **remark policies for access port, network port and hybrid port egress marking**

Remark policies are applied to access ports/SAPs, network ports/IP interfaces, and hybrid ports/SAPs/IP interfaces to configure the marking values for different forwarding classes and profiles. These policies enable marking of packet header QoS fields for packets that are forwarded out of the port, SAP, or IP interface. See [Remark policies](#) for more information.

- **queue management policies for buffer allocation and slope configuration on service egress and network port egress**

Queue management policies are applied to service egress, access port egress, network port egress, and hybrid port egress to configure CBS and MBS parameters for the egress queues and the WRED slope parameters for the queues.

The 7210 SAS-Mxp provides an option to use port-based queuing on access ports. This is a per-node configuration option and is mutually exclusive to the use of SAP-based egress queues (configured through service egress policies). When enabled, all SAPs on the access port share a set of 8 (eight) queues configured on the port and the access egress policy is used to define the queue parameters for port-based queues.

Service ingress, service egress, access ingress, access egress, and network QoS policies are defined with a scope of either **template** or **exclusive**. Template policies can be applied to multiple entities (such as SAPs and ports); exclusive policies can only be applied to a single entity.

The following policies are supported:

- one service ingress QoS policy and one service egress QoS policy applied to a specific SAP
- one access ingress and one access egress QoS policy applied to an access port
- one network QoS policy applied to a specific IP interface or network port, based on the type of network QoS policy; a network QoS policy defines both ingress and egress behavior
- one network queue policy applied to the network port

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

The 7210 SAS-Mxp can operate in either the **low** SAP scale mode or **high** SAP scale mode. In **low** SAP scale mode, SAP and service scaling is limited by the amount of CAM resources available for the SAP-ingress policy (both classification and meters). In the **high** SAP scale mode, SAP and service scaling is significantly higher compared to the **low** SAP scale mode and use table-based classification with ingress service meters. The use of network port policies remains unchanged when the system is operating in the **high** SAP scale mode.

SAPs configured on ports operating in hybrid mode cannot be configured to use access ingress QoS policies. Therefore, the **access-ingress-qos port-mode** option is not supported for ports configured in hybrid mode.

The following QoS policies are supported on access ports and SAPs in the low SAP scale mode:

- service ingress policy for SAP ingress classification and metering using the following:
 - CAM-based classification and metering/policing
 - table-based classification and CAM-based metering/policing
 - table-based classification and service-meter pool, also called table-based meter/policer pool, for metering/policing
- service egress policy for SAP egress queuing, shaping, and scheduling with an egress policy for marking only
- access egress policy for access port egress queuing, shaping, scheduling, and marking (this policy is mutually exclusive with the use of service egress policies)

The following QoS policies are supported on access ports and SAPs in the high SAP scale mode:

- Supports a choice of an access port ingress policy on access service delivery ports or SAP ingress policies on access ports. Nokia recommends using an access port ingress policy for higher SAP scaling.
- If using a service ingress policy for SAP ingress classification and metering, the following QoS policies are recommended:
 - Epipe and VPLS SAPs using ingress table-based classification and table-based policing on service delivery ports for higher SAP scaling
 - IES and VPRN SAPs using table-based classification or CAM-based classification
 - R-VPLS SAPs using CAM-based classification and policing
- If using an access ingress policy for access port ingress classification and metering, the following QoS policies are recommended:
 - access ingress policy for access port ingress to classify and police traffic
 - the user has an option to use a single policy per access port or share a single access port ingress policy with multiple access ports to reduce resource consumption
- Access egress policy for access port egress queuing, shaping, scheduling, and marking. This policy is mutually exclusive with the use of service egress policies.

A summary of the major functions performed by the QoS policies is listed in the following table.



Note:

Not all policies are supported on all platforms. See the following sections and chapters for more information.

Table 7: QoS policy types and descriptions for 7210 SAS-Mxp devices in network mode

Policy type	Applied at...	Description	Section
Service Ingress	SAP ingress	<ul style="list-style-type: none"> • defines up to 32 FC meters and meter parameters for traffic classification 	Service ingress QoS policies

Policy type	Applied at...	Description	Section
		<ul style="list-style-type: none"> defines match criteria to map flows to the meters based on any one of the criteria (optionally) can map the priority bits (IP DSCP (for all IP packets) and dot1p bits for tagged non-IP packets) to FC and profile for table-based classification (optionally) can use meters either from the CAM-based meter pool or the table-based service-meter pool 	
Service Egress	SAP Egress	<ul style="list-style-type: none"> defines up to 8 FC queues maps FCs to the queues defines Queue parameters (for example, rate, priority, weight, and so on) for the queue that determine how the queue gets the available bandwidth and prioritized by the scheduler defines FC to remarking values 	Service egress QoS policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12
Access Ingress	Access Port	<ul style="list-style-type: none"> provides an option to use a single access ingress QoS policy per port or share an access ingress QoS policy with multiple access ports defines dot1p and IP DSCP classification criteria to use; available with non-shared access-ingress policies defines IP criteria, dot1p, and IP DSCP classification criteria to use; available only with shared access-ingress QoS policies defines up to 16 meters per access ingress QoS policy 	Access ingress QoS policies
Access Egress	Access Port	<ul style="list-style-type: none"> defines FC to remarking values when port-based queuing is enabled, it is used to configure the queue parameters for port-based queues 	Access egress QoS policies on 7210 SAS-Mxp
Network (of type 'ip-interface')	IP interface	<ul style="list-style-type: none"> used for classification/marketing of MPLS packets at ingress, defines MPLS LSP-EXP to FC mapping and 16 meters used by FCs at egress, defines FC to MPLS LSP-EXP marking 	Network QoS policies in network mode
Network (of type 'port')	Network Ports and Hybrid Ports	<ul style="list-style-type: none"> used for classification/marketing of IP packets at ingress, defines DSCP or dot1p to FC mapping and 16 meters 	Network queue policies in network mode

Policy type	Applied at...	Description	Section
		<ul style="list-style-type: none"> at egress, defines FC to DSCP or dot1p marking or both 	
Network Queue	Network Ports and Hybrid Ports	<ul style="list-style-type: none"> defines FC mappings to network queues and queue characteristics for the queues 	Network queue policies in network mode
Slope	Access ports, Network ports and Hybrid ports	<ul style="list-style-type: none"> enables or disables the high-slope, low-slope, and non-TCP parameters within the egress pool 	Slope policy parameters
Remark	Network Port, Access ports and Hybrid Ports	<ul style="list-style-type: none"> applied at egress, defines the FC to Priority Bits (DSCP or dot1p or EXP) marking 	Remark policies on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (network mode)

2.1.4 Overview of QoS policies on 7210 SAS-R6 and 7210 SAS-R12

7210 SAS-R6 and 7210 SAS-R12 QoS policies are applied on service ingress, service egress, access port egress, network port ingress and egress, and network IP interface ingress. These policies allow users to configure the following:

- classification rules to map traffic to FCs
- FC association with meters and meter parameters used for policing (rate-limiting)
- queuing parameters for shaping, scheduling, and buffer allocation
- QoS marking and interpretation

The 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c support the following types of QoS policies:

- service ingress policies for access SAP ingress**

Service ingress QoS policies are applied to customer-facing SAPs. Traffic that enters through the SAP is classified to map to an FC. FCs are associated with meters or policers on ingress. Mapping traffic meters can be based on combinations of customer QoS marking (IEEE 802.1p bits), IP criteria, and MAC criteria. The policy defines the number of FC meters for unicast traffic and other meter characteristics (such as CIR, PIR, and so on). Each FC can be associated with different unicast parameters.

A service ingress QoS policy also defines up to three meters per FC for multipoint traffic for use with multipoint services. The system supports the configuration of up to 32 meters per service ingress QoS policy. In the case of VPLS, four forwarding types (not to be confused with forward classes) are supported: unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are flooded to all destinations within the service, while the unicast forwarding type is handled in a point-to-point fashion within the service.

Service ingress QoS policies are supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c.

- **service egress policies for access SAP egress**

Service egress QoS policies are applied to SAPs and map FCs to service egress queues for a service. The system allocates a maximum of eight queues per SAP for the eight FCs; this is not user configurable. All traffic types (unicast and BUM) share the same service egress queue. A service egress QoS policy defines the FC queue characteristics and how the FC to priority bits in the packet header are remarked (for example, IEEE 802.1p bits in the Ethernet VLAN tag) in the customer traffic.

On the 7210 SAS-R6 and 7210 SAS-R12, the user has a per-node or per-chassis option of configuring the SAP-based or access port-based egress queuing mode. The SAP-based egress queuing mode (**port-scheduler-mode** command disabled) uses a service egress QoS policy with the capability to use eight egress queues per SAP. The access port-based egress queuing mode (**port-scheduler-mode** command enabled) uses an access egress policy with the capability to use eight egress queues for all SAPs configured on the access port.

A service egress QoS policy or access egress policy also defines how to remark the FC-to-IEEE 802.1p bits in the customer traffic.

Service egress policies are supported only on 7210 SAS-R6 IMM-b and 7210 SAS-R12 IMM-b. 7210 SAS-R6 IMM-c and 7210 SAS-R12 IMM-c support only the port-based egress queuing mode and use access egress policies to configure the access port egress queues.

- **access egress policies for access port egress**

An access egress policy is applied to all SAPs on the physical port (and not the logical port (SAP)) for SAP-egress policies. It applies to all the SAPs on the port. Access egress policies provide different capabilities based on the egress queuing mode applied to the node. In SAP-based egress queuing mode, an access egress policy defines the remarking of the FC-to-packet header bits. For example, IEEE 802.1p bits in the Layer 2 VLAN header.

In port-based egress queuing mode, the access egress policy, in addition to remarking, is used to define the queuing and scheduling behavior for the port-based egress queues. Access egress QoS policies are applied to ports and map forwarding classes (FCs) to port egress queues. The system allocates a maximum of eight queues per port for the eight FCs. The allocation is not user-configurable. All traffic types, (unicast and BUM traffic types) share the same queue on port egress.

An access egress QoS policy defines the FC queue characteristics and the remarking of the FC to priority bits in the packet header (for example, IEEE 802.1p bits in the Ethernet VLAN tag) in the customer traffic.

Access egress policies are supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c. On IMM-b cards, when SAP-based egress queuing is configured, an access egress policy is used to configure marking. When port-based egress queuing is configured, an access egress policy is used to define port egress queue shaping and scheduling parameters, and configure marking.

- **access ingress policies for access port ingress**

An access ingress policy is applied to the physical port instead of the SAP; the policy applies to all SAPs configured on the specific access port. At ingress, the access ingress QoS policy uses dot1p, DEI with dot1p, or IP DSCP values to assign an FC and profile to traffic, which facilitates the classification of traffic received on the access port. The FC is associated with meters at ingress. An access ingress QoS policy allows the user to define up to one meter per FC for unicast traffic, and up to one meter per FC for multipoint traffic (broadcast, multicast, and unknown-unicast) for multipoint services. The system supports up to 16 meters per access ingress QoS policy.

This policy is supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c.



Note:

Access ingress policies are available only when the node is operating in **sap-scale mode high**. See the *7210 SAS-Mxp, S, Sx, T Services Guide* and *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information about this command.

- **network policies for network port and hybrid port ingress and egress, and network IP interface ingress**

For devices configured to operate in network mode, two types of network QoS policies are supported: one that is applied to a network IP interface and the other to a network port. On ingress, the policy applied to an IP interface maps incoming MPLS LSP EXP values to FC and profile state for traffic received from the core network. On egress, the network policy maps FC and profile state to MPLS LSP EXP values for traffic transmitted into the core network. The network policy applied to a network port maps incoming IP packets, DSCP, or dot1p values to the FC and the profile state for the traffic received from the core network. On egress, the network policy maps FC and profile state to DSCP or dot1p values for IP traffic transmitted into the core network.

These network policies are supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c.

- **network queue policies for network port and hybrid port, egress**

Network queue policies are applied on egress to network ports when operating in network mode. The policies define the FC queue characteristics for these entities. The FCs are mapped to the queues. The FC-to-queue mapping is static and not configurable. The number of queues is not user-configurable and the software allocates eight queues at the port level.

These network queue policies are supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c.

- **remark policies for service egress, access port egress, network port and hybrid port egress, and network IP interface egress marking**

Remark policies are applied to access ports/SAPs, network ports/IP interfaces, and hybrid ports/SAPs/IP interfaces to configure the marking values for different forwarding classes and profiles. These policies enable marking of packet header QoS fields for packets that are forwarded out of the port, SAP, or IP interface. See [Remark policies](#) for more information.

- **queue management policies for buffer allocation and slope configuration on service egress and network port egress**

Queue management policies are applied to service egress, access port egress, network port egress, and hybrid port egress to configure CBS and MBS parameters for the egress queues and the WRED slope parameters for the queues.

This policy is supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c. On the 7210 SAS-R6 IMM-c and 7210 SAS-R12 IMM-c, the CBS and MBS parameters are system-defined and not user-configurable. The values configured in the queue management policy for CBS and MBS parameters are ignored and only the WRED slope parameters are used.

Service ingress, access ingress, service egress, access egress, and network QoS policies are defined with a scope of either **template** or **exclusive**. Template policies can be applied to multiple entities (such as SAPs and ports); exclusive policies can only be applied to a single entity.

One service ingress and one service egress QoS policy can be applied to a specific SAP access egress policy, which can then be applied to an access port. One network QoS policy can be applied to a specific IP interface or network port based on the type of network QoS policy. A network QoS policy defines both ingress and egress behavior. One network queue policy can be applied to the network port.

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

The 7210 SAS-R6 and 7210 SAS-R12 can operate in either the **low** SAP scale mode or **high** SAP scale mode. In **low** SAP scale mode, SAP and service scaling is limited by the amount of CAM resources available for the SAP ingress policy (both classification and meters). In the **high** SAP scale mode, SAP and service scaling are significantly higher compared to the **low** SAP scale mode and use access port ingress policies. The use of network port policies remains unchanged when the system is operating in **high** SAP scale mode; however, the **high** SAP scale mode assumes that the user requires Layer 2 uplinks, and uses access port ingress and egress policies on those uplinks.

SAPs configured on ports operating in **hybrid** mode cannot be configured to use access ingress QoS policies. Therefore, the **access-ingress-qos port-mode** option is not supported for ports configured in **hybrid** mode.

The following QoS policies are supported on access ports and SAPs in the **low** SAP scale mode:

- service ingress policy for SAP ingress classification and metering using the following:
 - CAM-based classification and metering
 - table-based classification and CAM-based metering
- service egress policy for SAP egress queuing, shaping, and scheduling with an egress policy for marking only
- access egress policy for access port egress queuing, shaping, scheduling, and marking (this policy is mutually exclusive with the use of service egress policies)

The following QoS policies are supported on access ports and SAPs in the **high** SAP scale mode:

- the choice of an access port ingress policy on access service delivery ports or per-SAP ingress policies; Nokia recommends using an access port ingress policy for higher SAP scaling
- access egress policy for port egress queuing, shaping, scheduling, and marking (this policy is mutually exclusive with the use of service egress policies)

The following table describes the major functions performed by QoS policies for 7210 SAS-R6 and 7210 SAS-R12.

Table 8: QoS policy types and descriptions for 7210 SAS-R6 and 7210 SAS-R12

Policy type	Applied at	Description	Section
Service ingress	Access SAP ingress	<ul style="list-style-type: none"> • defines up to 32 FC meters and meter parameters for traffic classification • defines match criteria to map flows to the meters based on any one of the criteria (IP or MAC, or both IP and MAC) 	Service ingress QoS policies

Policy type	Applied at	Description	Section
		<ul style="list-style-type: none"> optionally, can map the priority bits (DSCP) to FC and profile using table-based classification supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c 	
Service egress	SAP egress	<ul style="list-style-type: none"> available only in SAP-based egress queuing mode defines up to eight FC queues Maps FCs to queues defines queue parameters for the queues defines FC-to-remarking values defines CIR levels and PIR weights that determine how the queue is prioritized by the scheduler supported on the 7210 SAS-R6 IMM-b and 7210 SAS-R12 IMM-b 	Service egress QoS policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12
Access Ingress	Access Port	<ul style="list-style-type: none"> defines dot1p and IP DSCP classification criteria to use defines up to 16 meters per access ingress QoS policy supported on 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 	Access ingress QoS policies

Policy type	Applied at	Description	Section
		7210 SAS-R12 IMM-c	
Access egress	Access port	<p>In port-based egress queuing mode, the access egress QoS policy supports the following:</p> <ul style="list-style-type: none"> • defines up to eight FC queues • maps FCs to queues • defines queue parameters for the queues • defines FC-to-remarking values • defines queue priorities that determine how the queue is prioritized by the scheduler • supported on both IMM-b and IMM-c cards <p>In SAP-based egress queuing mode, access egress is used to configure the FC-to-remarking map values</p>	Access egress QoS policies on 7210 SAS-R6 and 7210 SAS-R12
Network (type = ip-interface)	Network IP interface	<ul style="list-style-type: none"> • used for classification and marking of MPLS packets • on ingress, defines MPLS LSP-EXP-to-FC mapping, and defines up to 16 FC meters • on egress, defines FC-to-MPLS LSP-EXP marking • supported the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 	Network QoS policy "ip-interface" type

Policy type	Applied at	Description	Section
		7210 SAS-R12 IMM-c	
Egress rate	Access and network port	<ul style="list-style-type: none"> configures the maximum bandwidth available for traffic sent out of a specified port supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c 	Egress port rate limiting
Network (type = port)	Network and hybrid ports	<ul style="list-style-type: none"> used for classification and marking of IP packets on ingress, defines DSCP or dot1p-to-FC mapping, and up to 16 FC meters on egress, defines FC-to-DSCP or dot1p marking, or both supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c 	Network QoS policy "port" type
Network queue	Network and hybrid ports	<ul style="list-style-type: none"> defines FC mappings to network queues and queue characteristics for the queues supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c 	Network queue policies in network mode

Policy type	Applied at	Description	Section
Queue management policies	Queues at service ingress, service egress, and network egress	<ul style="list-style-type: none"> defines the CBS and MBS parameters for the queues enables or disables the high-slope and low-slope parameters for the queues supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c; on the 7210 SAS-R6 IMM-c and 7210 SAS-R12 IMM-c, only WRED slope parameters are used 	Queue management policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12
Remark	SAP egress, network egress (port and IP interface)	<ul style="list-style-type: none"> defines the FC-to-remarking values supported on the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, and 7210 SAS-R12 IMM-c 	Remark policies on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (network mode)

2.1.5 Overview of QoS policies on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE

QoS policies are applied on service ingress, access port egress, network port ingress and egress, and network IP interfaces ingress when configured 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE operates with MPLS uplinks.

These policies allow users to configure the following:

- classification rules for how traffic is mapped to FCs
- FC association with meters and meter parameters used for policing (rate-limiting)
- queuing parameters for shaping
- QoS marking/interpretation

There are several types of QoS policies:

- **service ingress policies for access SAP ingress**

Service ingress QoS policies are applied to the customer-facing SAPs. Traffic that enters through the SAP is classified to map it to an FC. FCs are associated with meters/policers on ingress. The mapping

of traffic to meters can be based on combinations of customer QoS marking (IEEE 802.1p bits), IP criteria, and MAC criteria.

- **access egress policies for access port egress**

Access egress policies are analogous to SAP egress policies as defined in the 7750 SR series of products. The difference is the point of attachment. An access egress policy is applied on the physical port as opposed to the logical port (SAP) for SAP egress policy. It applies to all SAPs on a port. An access egress QoS policy maps the traffic egressing on customer facing ports into various queues and marks the traffic accordingly.

- **access ingress policies for access port ingress**

An access ingress policy is applied to the physical port instead of the SAP; the policy applies to all SAPs configured on the specific access port. At ingress, the access ingress QoS policy uses dot1p, DEI with dot1p, or IP DSCP values to assign an FC and profile to traffic, which facilitates the classification of traffic received on the access port. The FC is associated with meters at ingress. An access ingress QoS policy allows the user to define up to one meter per FC for unicast traffic, and up to one meter per FC for multipoint traffic (that is, broadcast, multicast, and unknown-unicast) for multipoint services. The system supports up to 16 meters per access ingress QoS policy.



Note:

This policy is available only when the node is operating in **sap-scale mode high**. See the *7210 SAS-Mxp, S, Sx, T Services Guide* and *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information about the **sap-scale-mode** command.

- **network policies for network and hybrid port, ingress and egress**

The 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE support two types of network QoS policies, one applied to a network IP interface and the other to a network port or a hybrid port. Network QoS policies are applied to IP interfaces. On ingress, the policy applied to an IP interface maps incoming MPLS LSP EXP values to FC and profile state for the traffic received from the core network. On egress, the policy maps FC and profile state to MPLS LSP EXP values for traffic to be transmitted into the core network. The network policy applied to a network port maps incoming IP packets, DSCP, or dot1p values, to the FC and the profile state for the traffic received from the core network. On egress, the policy maps FC and profile state to DSCP or dot1p values for IP traffic to be transmitted into the core network.

- **network queue policies for network and hybrid port, egress**

Network queue policies are applied on egress to network or hybrid ports when operating in network mode. The policies define the FC queue characteristics for these entities. The FCs are mapped to the queues. There are 16 queues at the port level. FC-to-queue mapping is static and not configurable. The number of queues is not user-configurable, and the software allocates 16 queues at the port level.

- **port scheduler policies for access port, network port, and hybrid port egress**

Port scheduler policies are applied on egress for access, network, and hybrid ports. These policies allow the user to define queue scheduling attributes, such as strict-priority queuing and weighted queuing.

- **slope policies**

Slope policies are applied to the egress queues on the access, network, and hybrid ports. These policies define the WRED congestion management attributes, such as drop probability and thresholds for high-profile and low-profile traffic.

- **remark policies for access port, network port, and hybrid port egress marking**

Remark policies are applied to access ports/SAPs, network ports/IP interfaces, and hybrid ports/SAPs/IP interfaces to configure the marking values for different forwarding classes and profiles. These policies enable marking of packet header QoS fields for packets that are forwarded out of the port, SAP, or IP interface. See [Remark policies](#) for more information.

The characteristics of the FC meters, including the number of FC meters for unicast traffic and the meter characteristics (like CIR, PIR, and so on) are defined within the policy. Each FC can be associated with different unicast parameters. A service ingress QoS policy also defines up to three (3) meters per FC to be used for multipoint traffic for multipoint services. Up to 32 meters, in total, are supported per Service ingress QoS policies.

In the case of the VPLS, the following types of forwarding are supported (which is not to be confused with FCs): 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 point-to-point in the service.

The FCs are mapped to 16 queues at the port level (8 for unicast and 8 for multicast). FC-to-queue mapping is static and not configurable. The number of queues is not user configurable and the software allocates 16 queues at the port level. An access egress policy also defines the remarking of the FC-to-packet header bits (for example, IEEE 802.1p bits in the Layer 2 VLAN header, and others.).

The following applies to the use of QoS policies on hybrid ports:

- Network queue policies are supported for queue configuration of egress queues on hybrid ports. These egress queues are shared by traffic sent out of SAPs and network IP interfaces configured on hybrid ports.
- Network QoS policies of type "ip-interface" are supported for network IP interfaces on hybrid ports. The behavior is similar to the existing behavior for network IP interfaces on network ports. It supports per IP interface ingress classification and policing and egress marking (only EXP marking for MPLS traffic).
- Network QoS (type = port) policies are supported for hybrid ports. The behavior is similar to existing behavior for network ports. It supports per port ingress classification and policing and egress marking (dot1p and/or DSCP marking) for IP control packets.
- SAP ingress QoS policies are supported for SAPs configured on hybrid ports. The behavior is similar to existing behavior for access SAP ingress. It supports per SAP ingress classification and policing.
- For marking traffic sent out of SAPs and IP traffic sent out of IP interfaces configured on hybrid ports, users must use the network QoS policy of type "port", with an option to mark dot1p, DSCP, or both.



Note:

If DSCP remarking or both is specified, the DSCP field is not marked for the traffic sent out of the Layer 2 SAPs.

Service ingress, access egress, and network QoS policies are defined with a scope of either **template** or **exclusive**. Template policies can be applied to multiple entities (such as SAPs and ports); exclusive policies can be applied to only a single entity. One service ingress QoS policy can be applied to a specific SAP.

An access ingress policy is applied to the physical port instead of the SAP; the policy applies to all SAPs configured on the specific access port. At ingress, the access ingress QoS policy uses dot1p, DEI with dot1p, or IP DSCP values to assign a forwarding class and profile to traffic, which facilitates the classification of traffic received on the access port. The FC is associated with meters at ingress. An access ingress QoS policy allows the user to define up to one meter per forwarding class for unicast traffic, and up to one meter per forwarding class for multipoint traffic (that is, broadcast, multicast, and unknown-unicast) for multipoint services. The system supports up to 16 meters per access ingress QoS policy.

An access egress policy can be applied to an access port. One access egress QoS policy can be applied to the access port. One network QoS policy can be applied to a specific IP interface, network port, or hybrid port based on the type of network QoS policy. A network QoS policy defines both ingress and egress behavior. One network queue policy can be applied to the network port or hybrid port. If no QoS policy is applied to a SAP, port, or interface, a default QoS policy is applied.

The 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE can operate in either the **low** SAP scale mode or **high** SAP scale mode. In **low** SAP scale mode, SAP and service scaling is limited by the amount of CAM resources available for the SAP ingress policy (both classification and meters). In the **high** SAP scale mode, SAP and service scaling are significantly higher compared to the **low** SAP scale mode and use access port ingress policies. The use of network port policies remains unchanged when the system is operating in **high** SAP scale mode; however, the **high** SAP scale mode assumes that the user requires Layer 2 uplinks, and uses access port ingress and egress policies on those uplinks.

SAPs configured on ports operating in **hybrid** mode cannot be configured to use access ingress QoS policies. Therefore, the **access-ingress-qos port-mode** option is not supported for ports configured in **hybrid** mode.

The following QoS policies are supported on access ports and SAPs in the **low** SAP scale mode:

- service ingress policy for SAP ingress classification and metering using the following:
 - CAM-based classification and metering
 - table-based classification and CAM-based metering
- service egress policy for SAP egress queuing, shaping, and scheduling with an egress policy for marking only
- access egress policy for access port egress queuing, shaping, scheduling, and marking (this policy is mutually exclusive with the use of service egress policies)

The following QoS policies are supported on access ports and SAPs in the **high** SAP scale mode:

- the choice of an access port ingress policy on access service delivery ports or per-SAP ingress policies; Nokia recommends using an access port ingress policy for higher SAP scaling
- access egress policy for port egress queuing, shaping, scheduling, and marking (this policy is mutually exclusive with the use of service egress policies)

The following table describes the major functions performed by QoS policies.

Table 9: QoS policy types and descriptions for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE devices in network mode


Policy type	Applied at...	Description	Section
Service ingress	SAP ingress	<ul style="list-style-type: none"> • defines up to 16 FC meters and meter parameters for traffic classification • defines match criteria to map flows to the meters, based on any one of the criteria 	Service ingress QoS policies
Access egress	Access port	<ul style="list-style-type: none"> • defines up to 8 FC queues and queue parameters for traffic classification • maps FCs to the queues • defines FC to remarking values 	Access egress QoS policies on 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE

Policy type	Applied at...	Description	Section
		<ul style="list-style-type: none"> defines CIR levels and PIR weights that determine how the queue is prioritized by the scheduler 	
Access Ingress	Access Port	<ul style="list-style-type: none"> defines dot1p and IP DSCP classification criteria to use defines up to 16 meters per access ingress QoS policy 	Access ingress QoS policies
Network (of type "ip-interface")	IP interface	<ul style="list-style-type: none"> used for classification/marketing of MPLS packets at ingress, defines MPLS LSP-EXP to FC mapping and 12 meters used by FCs at egress, defines FC to MPLS LSP-EXP marking 	Network QoS policies in network mode
Network (of type "port")	Network and hybrid ports	<ul style="list-style-type: none"> used for classification/marketing of IP packets at ingress, defines DSCP or dot1p to FC mapping and 8 meters at egress, defines FC to DSCP, dot1p marking, or both 	Network QoS policies in network mode
Network queue	Network ports and hybrid ports	<ul style="list-style-type: none"> defines FC mappings to network queues and queue characteristics for the queues 	Network queue policies in network mode
Slope	Access ports, network ports, and hybrid ports	<ul style="list-style-type: none"> enables or disables the high-slope, low-slope, and non-TCP parameters within the egress pool 	Slope policy parameters
Port scheduler	Access ports, network ports, and hybrid ports	<ul style="list-style-type: none"> defines the parameters for the port scheduler 	Port scheduler policies for 7210 SAS-T
Remark	Network port, access ports, and hybrid ports	<ul style="list-style-type: none"> applied at egress, defines the FC-to-priority bits (DSCP or dot1p or EXP) marking 	Remark policies on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (network mode)

2.2 Network and service QoS policies

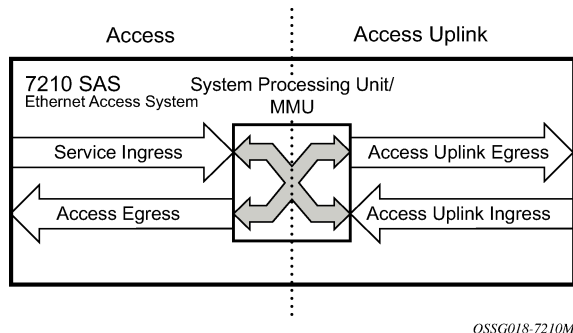
The QoS mechanism within the 7210 SAS is specialized for the type of traffic on the interface. For customer interfaces, service ingress and access service egress traffic exists, and for IP interfaces, network ingress and network egress traffic exists (as shown in the following figure).

Figure 1: 7210 SAS traffic types operating in network mode



When operating in access-uplink mode, the QoS mechanisms available are similar to network mode, except that network ingress and network egress traffic is associated with access-uplink interfaces instead of network IP interface or network ports (as shown in the following figure).

Figure 2: 7210 SAS-T traffic types for access-uplink mode



The 7210 SAS uses QoS policies applied to a SAP, a network port, access port, or access-uplink port to define queuing, queue attributes, meter/policer attributes and QoS marking interpretation.

The 7210 SAS supports the following types of network and service QoS policies: [Network QoS policies in network mode](#), [Network queue QoS policies](#), [Service ingress QoS policies](#), and [Service egress QoS policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#).

2.2.1 Network QoS policies in network mode

The following applies to QoS policies configured in network mode.

- The following types of network QoS policies can be defined: **ip-interface** and **port**. By default, when a network QoS policy is created, it is of the **ip-interface** type.
- Create a network QoS policy of the **ip-interface** type using the **configure>qos>network** context.

- Create a network QoS policy of the **port** type using the **configure>qos>network** context.
- When a network QoS policy of the **ip-interface** type is applied to an IP interface configured on network port and hybrid ports, the policy is used for classification of MPLS packets received based on LSP-EXP bits and marking of MPLS-EXP bits for MPLS traffic sent out of the IP interface.
- When a network QoS policy of the **port** type is applied to a network and hybrid port, it is used for classification of IP packets, based on the DSCP or dot1p bits and marking of DSCP or dot1p bits for packets sent out of network or hybrid ports.
- On 7210 SAS-R6 and 7210 SAS-R12, an option is available to reallocate resources needed by network QoS ingress to other features sharing the ingress-internal-tcam resource pool. Users must ensure that sufficient resources are available for network ingress QoS, if the user intends to use network ports and network IP interfaces.

2.2.1.1 Network QoS policy “ip-interface” type

Network QoS policies of the **ip-interface** type define ingress FC meters and map traffic to those meters for IP interfaces. When a network QoS policy is created, it always has two meters defined that cannot be deleted: one for the unicast traffic and one for the multipoint traffic. These meters exist within the definition of the policy. The meters are used by the hardware only when the policy is applied to an IP interface. This policy also defines the FC to EXP bit marking, on the egress mode.

A network QoS policy defines both the ingress and egress handling of QoS on the network IP interface and network port. The following functions are defined for a network policy of the **ip-interface** type:

- Ingress
 - defines EXP value mapping to FCs. Ingress profile assignment using MPLS EXP values is configured using the `mpls-lsp-exp-profile-map` policy.
 - defines FC to meter mapping. By default, meters are color aware. The user cannot disable the meters or change the meter color mode (the meter color mode is always set to color-aware).
- Egress
 - defines the FC to EXP value markings
 - remarking of QoS bits can be enabled or disabled



Note:

See [Remark policies](#) for more information about MPLS EXP, IP DSCP, and dot1p marking using network QoS policies.

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

- a unique network QoS policy ID
- egress FC-to-EXP value mappings for each FC
- a default ingress FC and in-profile/out-of-profile state
- at least one default unicast FC meter. See [Meter/policer parameters](#) for information about the parameters that can be configured for a meter.
- optional multipoint FC meter

Optional network QoS policy elements include:

- additional unicast meters up to a total of 8

- additional multipoint meters up to 8
- EXP value to FC and profile state mappings for all EXP values received

Network policy ID 2 is reserved as the default network QoS policy of the **ip-interface** type. The default policy cannot be deleted or changed.

Default network QoS policy 2 is applied to all IP interfaces that do not have another network QoS policy assigned.

The network QoS policy applied at network egress (for example, on an IP interface) determines how or whether the profile state is marked in packets transmitted to the service core network. If the profile state is marked in the service core packets, out-of-profile packets are preferentially dropped over in-profile packets at congestion points in the core network. For network egress, traffic remarking in the network QoS policy is disabled. The following table lists the default mapping of FC-to-EXP values.

Table 10: Default network QoS policy (type = ip-interface) egress marking

FC-ID	FC name	FC label	Egress EXP marking	
			In-profile	Out-of-profile
7	Network Control	nc	111 - 7	111 - 7
6	High-1	h1	110 - 6	110 - 6
5	Expedited	ef	101 - 5	101 - 5
4	High-2	h2	100 - 4	100 - 4
3	Low-1	l1	011 - 3	010-2
2	Assured	af	011-3	010 - 2
1	Low-2	l2	001 - 1	001 - 1
0	Best Effort	be	000 - 0	000 - 0

For network ingress, the following table lists the default mapping of EXP values to FC and profile state for the default network QoS policy. Color-aware policing is supported on network ingress.

Table 11: Default network QoS policy (type = ip-interface) EXP-to-FC mapping

EXP value	FC ingress	Profile
0	be	Out
1	l2	In
2	af	Out
3	af	In
4	h2	In
5	ef	In

EXP value	FC ingress	Profile
6	h1	In
7	nc	In

2.2.1.2 Network QoS policy “port” type

Network QoS policy of the **port** type defines ingress FC meters and maps traffic to the meters for only IP traffic received on network and hybrid ports. When a network policy of this type is created, it has a single unicast meter that cannot be deleted. These meters exist within the definition of the policy. The meters are instantiated in hardware only when the policy is applied to a network port. This policy also defines the FC-to-DSCP or dot1p marking used for packets sent out through that port.

A network QoS policy of the **port** type defines both the ingress and egress handling of QoS on the network port.

The following functions are defined:

- Ingress
 - defines DSCP or dot1p value mapping to FCs. Only one type is supported, such as DSCP or dot1p, per policy, with an option to use the DEI bit along with dot1p classification for profile assignment.
 - defines FC-to-meter mapping. By default, meters are color aware. The user cannot disable meters or change the meter color mode (meter color mode is always set to **color-aware**).
- Egress
 - specifies the remark policy that defines FC-to-DSCP or dot1p (or both) value markings. See [Remark policies](#) for more information.

The following are the required elements defined in a network QoS policy of the **port** type:

- a unique network QoS policy ID and network-policy-type set to **port**
- egress FC-to-DSCP or dot1p (or both) value mappings for each FC
- a default ingress FC and in-profile/out-of-profile state
- at least one default unicast FC meter. See [Meter/policer parameters](#) for information about the parameters that can be configured for a meter.

Optional network QoS policy elements include the following:

- additional unicast meters up to a total of 8
- additional multipoint meter up to a total of 8
- DSCP or dot1p (or both) value to FC and profile state mappings for all DSCP or dot1p values received
- option to use the DEI bit along with dot1p classification for profile state mapping

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

The default network QoS policy is applied to all network ports that do not have another network QoS policy assigned.

The following table lists the default mapping of FC-to-dot1p and DSCP values.

Table 12: Default network QoS policy of type "port" egress marking

FC-ID	FC name	FC label	Egress DSCP marking		Egress dot1p marking	
			In-profile	Out-of-profile	In-profile	Out-of-profile
7	Network Control	nc	nc2	nc2	111 - 7	111 - 7
6	High-1	h1	nc1	nc1	110-6	110-6
5	Expedited	ef	ef	ef	101-5	101-5
4	High-2	h2	af41	af41	100-4	100-4
3	Low-1	l1	af21	af22	011-3	010-2
2	Assured	af	af11	af12	011-3	010-2
1	Low-2	l2	cs1	cs1	001-1	001-1
0	Best Effort	be	be	be	000-0	000-0

The following table lists the default mapping of dot1p or DSCP values to FC and profile state for the default network QoS policy of the **port** type for network ingress. Color-aware policing is supported on network ingress.

Table 13: Default network QoS policy of type "port": dot1p/DSCP-to-FC mapping

DSCP value	Dot1p value	FC ingress	Profile
be	0	be	Out
cs1	1	l2	In
af12	2	af	Out
af11	3	af	In
af41	4	h2	In
ef	5	ef	In
nc1	6	h1	In
nc2	7	nc	In

2.2.2 Network QoS policies in access-uplink mode

Network QoS policies define ingress FC meters and map traffic to the meters for access-uplink ports. A network QoS policy always has two meters/policers defined that cannot be deleted, one for the unicast traffic and one for multipoint traffic. These meters exist within the definition of the policy. The meters are

instantiated in hardware only when the policy is applied to an access-uplink port. The policy also defines the FC-to-priority bit marking, on egress.

A network QoS policy defines both the ingress and egress handling of QoS on the access-uplink ports. The following functions are defined:

- Ingress
 - defines dot1p value mapping to FCs and profiles with an option to use the DEI bit along with dot1p classification (DSCP is not available for use)
 - defines FC to meter mapping
- Egress
 - option to define the FC to dot1p value and IP DSCP value for marking
 - remarking of QoS bits can be enabled or disabled

The following are the required elements defined in a network QoS policy:

- unique network QoS policy ID
- egress FC to dot1p value mappings for each FC
- default ingress FC and in-profile/out-of-profile state
- at least one default unicast FC meter. See [Meter/policer parameters](#) for information about the parameters that can be configured for a meter.
- at least one multipoint FC meter

Optional network QoS policy elements include the following:

- additional unicast meters up to a total of 8
- additional multipoint meters up to 8
- dot1p value to FC and profile state mappings for all dot1p values received
- option to use the DEI bit along with dot1p classification for profile state mapping

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

The default network QoS policy is applied to all access-uplink ports that do not have another network QoS policy assigned. The network QoS policy applied at network egress (for example, on an access-uplink port) determines how or whether the profile state is marked in packets transmitted into the service core network.

If the profile state is marked in the service core packets, out-of-profile packets are preferentially dropped over in-profile packets at congestion points in the core network. For network egress, traffic remarking in the network QoS policy is always enabled. The following table lists the default mapping of FC-to-dot1p values.

Table 14: Default network QoS policy used for egress marking on access-uplink ports

FC-ID	FC Name	FC Label	DiffServ Name	Egress dot1p marking	
				In-profile	Out-of-profile
7	Network Control	nc	NC2	111-7	111-7
6	High-1	h1	NC1	110-6	110-6
5	Expedited	ef	EF	101-5	101-5
4	High-2	h2	AF4	100-4	100-4

FC-ID	FC Name	FC Label	DiffServ Name	Egress dot1p marking	
				In-profile	Out-of-profile
3	Low-1	l1	AF2	011-3	010-2
2	Assured	af	AF1	011-3	010-2
1	Low-2	l2	CS1	00-1	001-1
0	Best Effort	be	BE	000-0	000-0

For network ingress, the following table lists the default mapping of dot1p values to FC and profile state for the default network QoS policy. Color-aware policing is supported on ingress for access-uplink ports.

Table 15: Default network QoS policy used for dot1p to FC on access-uplink ports

Dot1p value	FC ingress	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

2.2.3 Network queue QoS policies

This section provides information about network queue QoS policies.

2.2.3.1 Network queue policies in network mode

In the network mode of operation, network queue policies define the network FC queue characteristics. Network queue policies are applied on egress on network and hybrid ports for 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T operating in network mode. The system allocates a fixed number of queues for the network port, and FCs are mapped to the queues. All policies use a fixed number of queues, like the default network queue policy.

The following is the number of queues allocated for a network queue policy:

- On 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T, 8 (eight) queues are allocated and 8 (eight) FCs are mapped to the 8 (eight) queues. [Table 31: Forwarding class to queue-ID map](#) lists the FC-to-queue mapping.

- On 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, 16 queues are allocated, with 2 queues per FC, one each for unicast traffic and multicast (BUM) traffic. [Table 32: Forwarding class-to-queue ID map for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#) lists the FC-to-queue mapping.

On 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T, the network queues on hybrid ports are used for MPLS, IP, and SAP traffic sent out of IP interfaces and SAPs configured on hybrid ports.

The following queue characteristics can be configured on a per-FC basis:

- peak information rate (PIR) as a percentage of egress port bandwidth
- committed information rate (CIR) as a percentage of egress port bandwidth
- committed burst size (CBS) (using the queue-management policies); supported only on the 7210 SAS-Mxp, 7210 SAS-R6 IMM-b, and 7210 SAS-R12 IMM-b
- maximum burst size (MBS) (using the queue-management policies); supported only on the 7210 SAS-Mxp, 7210 SAS-R6 IMM-b, and 7210 SAS-R12 IMM-b
- queue-mode, strict or weighted; supported only on the 7210 SAS-Mxp, 7210 SAS-R6 IMM-b, and 7210 SAS-R12 IMM-b
- adaptation rules for CIR and PIR
- WRED slope parameters

Network queue policies are identified with a unique policy name, which conforms to the standard 7210 SAS alphanumeric naming conventions. The system default network queue policy is named "default" and cannot be edited or deleted.



Note:

CBS and MBS values are system-defined and cannot be provisioned by the user on the 7210 SAS-R6 IMM-c, 7210 SAS-R12 IMM-c, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T.

The following table lists the default network queue policy definition for the 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T configured in network mode.

Table 16: Default network queue policy definition (for 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T configured in network mode)

Forwarding class	Queue	Definition
Network-Control (nc)	Queue 8	<ul style="list-style-type: none"> PIR = 100% CIR = 10% CBS = 12.5
High-1 (h1)	Queue 7	<ul style="list-style-type: none"> PIR = 100% CIR = 10% CBS = 12.5%
Expedited (ef)	Queue 6	<ul style="list-style-type: none"> PIR = 100% CIR = 100% CBS = 12.5%

Forwarding class	Queue	Definition
High-2 (h2)	Queue 5	<ul style="list-style-type: none"> PIR = 100% CIR = 100% CBS = 12.5%
Low-1 (l1)	Queue 4	<ul style="list-style-type: none"> PIR = 100% CIR = 25% CBS = 12.5%
Assured (af)	Queue 3	<ul style="list-style-type: none"> PIR = 100% CIR = 25% CBS = 12.5%
Low-2 (l2)	Queue 2	<ul style="list-style-type: none"> PIR = 100% CIR = 25% CBS = 12.5%
Best-Effort (be)	Queue 1	<ul style="list-style-type: none"> PIR = 100% CIR = 0% CBS = 12.5%

The following table lists the default network queue policy definition for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 configured in network mode.

Table 17: Default network queue policy definition (for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 configured in network mode)

Forwarding class	Queue	Definition
Network-Control (nc)	Queue 8	<ul style="list-style-type: none"> PIR = 100% CIR = 10% CBS = 12.5 queue-mgmt = default queue-mode = weighted weight = 1
High-1 (h1)	Queue 7	<ul style="list-style-type: none"> PIR = 100% CIR = 10% CBS = 12.5% queue-mgmt = default queue-mode = weighted weight = 1

Forwarding class	Queue	Definition
Expedited (ef)	Queue 6	<ul style="list-style-type: none"> PIR = 100% CIR = 100% CBS = 12.5% queue-mgmt = default queue-mode = weighted weight = 1
High-2 (h2)	Queue 5	<ul style="list-style-type: none"> PIR = 100% CIR = 100% CBS = 12.5% queue-mgmt = default queue-mode = weighted weight = 1
Low-1 (l1)	Queue 4	<ul style="list-style-type: none"> PIR = 100% CIR = 25% CBS = 12.5% queue-mgmt = default queue-mode = weighted weight = 1
Assured (af)	Queue 3	<ul style="list-style-type: none"> PIR = 100% CIR = 25% CBS = 12.5% queue-mgmt = default queue-mode = weighted weight = 1
Low-2 (l2)	Queue 2	<ul style="list-style-type: none"> PIR = 100% CIR = 25% CBS = 12.5% queue-mgmt = default queue-mode = weighted weight = 1
Best-Effort (be)	Queue 1	<ul style="list-style-type: none"> PIR = 100% CIR = 0% CBS = 12.5%

Forwarding class	Queue	Definition
		<ul style="list-style-type: none"> queue-mgmt = default queue-mode = weighted weight = 1

2.2.3.2 Network queue policies in access-uplink mode

In access-uplink mode of operation, network queue policies are applied at egress of access-uplink ports for 7210 SAS-T devices operating in access-uplink mode. The system allocates 8 (eight) queues for the network port and FCs are mapped to the 8 (eight) queues. All policies uses 8 (eight) queues, like the default network queue policy.

The following queue characteristics can be configured on a per-FC basis:

- Peak Information Rate (PIR) as a percentage of egress port bandwidth
- Committed Information Rate (CIR) as a percentage of egress port bandwidth
- adaptation rule for CIR and PIR
- WRED slope parameters

Network queue policies are identified with a unique policy name, which conforms to the standard 7210 SAS alphanumeric naming conventions. The system default network queue policy is named "default" and cannot be edited or deleted. CBS values cannot be provisioned. The following table lists the default network queue policy definition in access-uplink mode.

Table 18: Default network queue policy definition (for 7210 SAS-T configured in access-uplink mode)

Forwarding class	Queue	Definition
Network-Control (nc)	Queue 8	<ul style="list-style-type: none"> PIR = 100% CIR = 10% CBS = 7%
High-1 (h1)	Queue 7	<ul style="list-style-type: none"> PIR = 100% CIR = 10% CBS = 7%
Expedited (ef)	Queue 6	<ul style="list-style-type: none"> PIR = 100% CIR = 100% CBS = 21%
High-2 (h2)	Queue 5	<ul style="list-style-type: none"> PIR = 100% CIR = 100% CBS = 21%
Low-1 (l1)	Queue 4	<ul style="list-style-type: none"> PIR = 100% CIR = 25%

Forwarding class	Queue	Definition
		<ul style="list-style-type: none"> CBS = 7%
Assured (af)	Queue 3	<ul style="list-style-type: none"> PIR = 100% CIR = 25% CBS = 21%
Low-2 (l2)	Queue 2	<ul style="list-style-type: none"> PIR = 100% CIR = 25% CBS = 7%
Best-Effort (be)	Queue 1	<ul style="list-style-type: none"> PIR = 100% CIR = 0% CBS = 7%

2.2.4 Service ingress QoS policies

Service ingress QoS policies define ingress service FC meters and map flows to the meters. When a service ingress QoS policy is created, it has a single meter defined that cannot be deleted and is used for all the traffic (both unicast and multicast traffic). These meters exist within the definition of the policy, but are instantiated in hardware only when the policy is applied to a SAP. In the case where the service does not have multipoint traffic, the multipoint meters are not instantiated.

In the simplest service ingress QoS policy, all traffic is handled as a single flow and mapped to a single meter, including all flooded traffic. The required elements to define a service ingress QoS policy are the following:

- unique service ingress QoS policy ID
- QoS policy scope of **template** or **exclusive**
- number of classification and meter resources to allocate for this policy
- allocation of resources from the ingress internal CAM resource pool for use with service ingress QoS policies

Additionally, the allocation of resources to the appropriate classification match criteria.

- at least one default FC meter

See [Meter/policer parameters](#) for information about the parameters that can be configured for a meter.

Optional service ingress QoS policy elements include the following:

- additional unicast meters up to a total of 8 (eight)
- additional multipoint meters up to 32
- QoS policy match criteria to map packets to an FC
- option to use dot1p or IP DSCP table-based classification in Layer 2 services (Epipe and VPLS) on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12
- option to use IP DSCP table-based classification in Layer 3 services on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

- option to use meters from either CAM-based meter resource pool or table-based service-meter resource pool on 7210 SAS-Mxp

The following options are available when using resources for classification and policing:

- **CAM-based classification and policing**

Resources for both classification and policing are allocated from the CAM-based classification and meter resource pool. Configure resources from the CAM-based pool using the **configure>system>resource-profile>ingress-internal-tcam>qos-sap-ingress-resource** command. The user has the flexibility to use IP criteria (both IPv4 and IPv6) and MAC criteria for classification of traffic flows to an FC (FC). Using this option allows each SAP to define its FC-to-meter map, which is the default option that is enabled when the system boots up using the default configuration. This option is supported on the following platforms: 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (in network and access-uplink mode). This is available with all services (Epipe SAPs, VPLS SAPs, VPRN SAPs, IES SAPs, and R-VPLS SAPs).

- **table-based classification and CAM-based policing**

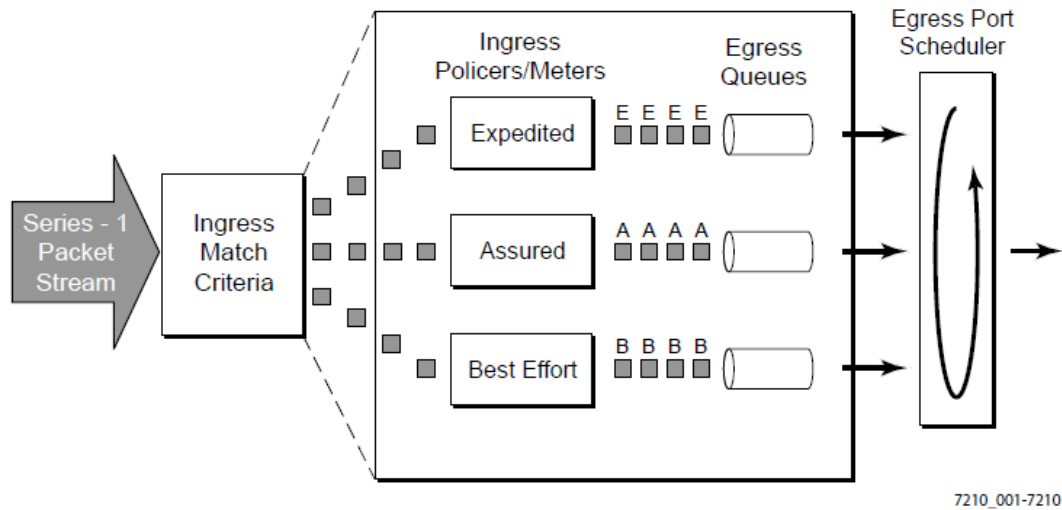
Resources for classification are allocated from the table-based classification pool of resources, and policing resources are allocated from the CAM-based meter resource pool. Configure resources for the CAM-based pool using the **configure>system>resource-profile>ingress-internal-tcam>qos-sap-ingress-resource** command. See [Service ingress QoS policies](#) for examples that show how to configure this option. Using this option, the classification of traffic flows can be done only using IP DSCP and or dot1p bits. (See [Service ingress QoS policies](#) for information about classification support details.) The use of IP and MAC criteria are mutually exclusive. Using this option allows each SAP to define its FC-to-meter map. This option is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12. This is available with all services (Epipe SAPs, VPLS SAPs, VPRN SAPs, IES SAPs, and R-VPLS SAPs).

- **table-based classification and service-meter based policing**

Resources for classification are allocated from the table-based classification pool of resources, and policing resources are allocated from the table-based service-meter resource pool. See [Service ingress QoS policies](#) for examples that show how to configure this option. With this option, the classification of traffic flows can be done only using IP DSCP or dot1p bits (for classification support details, see [Service ingress QoS policies](#)). The use of IP and MAC criteria is not available. All SAPs in the node must use a single FC-to-meter map (other than the default). This option is only supported on the 7210 SAS-Mxp and is only available for Epipe SAPs, VPLS SAPs, VPRN SAPs, and IES SAPs. It is not available for R-VPLS SAPs.

Each meter can have unique meter parameters to allow individual policing of the flow mapped to the FC. The following figure shows service traffic being classified into three different FCs.

Figure 3: Traffic queuing model for forwarding classes



Mapping a flow to an FC is controlled by comparing each packet to the match criteria in the QoS policy. The ingress packet classification to FC requires a provisioned classification policy.

2.2.4.1 CAM-based classification

When using CAM-based classification on 7210 SAS devices, at SAP ingress users have an option to use either MAC criteria or IP criteria, or both IPv4 and MAC criteria to allow users to use the available CAM classification resources effectively.

The following options are available:

- supported MAC header fields using the **mac-criteria any** option
- only dot1p bits in the MAC header using the **mac-criteria dot1p-only** option
- supported IPv4 header fields using the **ip-criteria any** option
- only IPv4 DSCP in the IPv4 header using the **ip-criteria dscp-only** option
- supported IPv6 header fields using the **ipv6-criteria any** option
- only IPv6 DSCP bits in the IPv6 header using the **ipv6-criteria dscp-only** option
- both MAC and IPv4 header fields using the both MAC and IPv4 criteria option together in a policy

Among the preceding supported criteria, the following can be configured in a single policy:

- **mac-criteria any**
- **mac-criteria dot1p-only**
- **ip-criteria any** and/or **ipv6-criteria any** or **ipv6-criteria dscp-only**
- **ip-criteria dscp-only** and/or **ipv6-criteria any** or **ipv6-criteria dscp-only**
- **mac-criteria any** and **ip-criteria any** or **ip-criteria dscp-only** and/or **ipv6-criteria dscp-only**
- **mac-criteria dot1p-only** and **ip-criteria any** or **ip-criteria dscp-only** and/or **ipv6-criteria dscp-only**

**Note:**

When specifying both MAC and IP criteria in a SAP ingress policy, only an IPv6 DSCP match is allowed. Other IPv6 fields, such as **src-address** and **dst-address**, are not allowed.

In addition to the preceding list of classification rules, the user can set the DEI bit for identifying the ingress profile and enabling color-aware policing. See [Discard eligibility indicator-based \(DEI-based\) classification and marking](#) and [Service ingress QoS policies](#) for more information. The packet fields that can be used as match criteria for SAP ingress classification are described in [Table 19: Service ingress QoS policy match criteria for 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T in network mode](#) and [Table 20: Service ingress QoS policy criteria for 7210 SAS-T in access-uplink mode](#).

**Note:**

To determine the resource allocation required for each of these different criteria, see [Service ingress QoS policies](#).

The IP and MAC match criteria can be very basic or quite detailed. IP and MAC match criteria are constructed using 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 an action that specifies the FC of packets that match the entry.

The entries are evaluated in numerical order based on the entry ID, from the lowest to highest ID value. The first entry that matches all match criteria has its action performed.

Table 19: Service ingress QoS policy match criteria for 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T in network mode

Match criteria	Description
IP criteria	IP DSCP value/mask and IP Precedence value (available for SAPs in VPLS, VLL, PBB Epipe I-SAP, PBB VPLS I-SAP, IES and VPRN, and R-VPLS services)
	IP source and mask, IP destination and mask, IP protocol, TCP/UDP source port, TCP/UDP destination port, (available only for SAPs in VPLS, VLL, PBB Epipe I-SAP, PBB VPLS I-SAP, IES and VPRN services)
IPv6 criteria	IP DSCP value/mask and IP Precedence value (available for SAPs in VPLS, VLL, PBB services)
	IPv6 128-bit source and mask, IPv6 128-bit destination and mask, IP protocol/next-header, TCP/UDP source port, TCP/UDP destination port, (available only for SAPs in VPLS, VLL, PBB Epipe I-SAP, PBB VPLS I-SAP)
MAC criteria	IEEE 802.1p/dot1p value/mask, Source MAC address/mask, Destination MAC address/mask, EtherType Value/Mask (available for VLL, VPLS, PBB (Epipe I-SAP, VPLS I-SAP, B-SAP), IES, VPRN, and R-VPLS services)

Table 20: Service ingress QoS policy criteria for 7210 SAS-T in access-uplink mode

Match criteria	Description
IP criteria	IP DSCP value/mask and IP Precedence value (available for access SAPs in VPLS, VLL, IES, and R-VPLS services)
	IP source and mask, IP destination and mask, IP protocol, TCP/UDP source port, TCP/UDP destination port, (available only for access SAPs in VPLS, VLL, and IES services)
IPv6 criteria	IP DSCP value/mask and IP Precedence value (available for SAPs in VPLS, and VLL services)
	IPv6 128-bit source and mask, IPv6 128-bit destination and mask, IP protocol/next-header, TCP/UDP source port, TCP/UDP destination port, (available only for SAPs in VPLS and VLL services)
MAC criteria	IEEE 802.1p/dot1p value/mask, Source MAC address/mask, Destination MAC address/mask, EtherType Value/Mask (available for VLL, VPLS, IES, and R-VPLS services)

The following table lists the MAC match criteria that can be used for an Ethernet frame, which depends on the frame format.

Table 21: MAC match Ethernet frame types

Frame format	Description
802.3	IEEE 802.3 Ethernet frame. Only the source MAC, destination MAC and IEEE 802.1p value are compared for match criteria.
Ethernet-II	Ethernet type II frame where the 802.3 length field is used as an Ethernet type (Etype) value Etype values are two byte values greater than 0x5FF (1535 decimal).

The following table lists the criteria that can be matched for the various MAC frame types.

Table 22: MAC match criteria frame type dependencies

Frame format	Source MAC	Dest MAC	IEEE 802.1p value	Etype value
802.3	Yes	Yes	Yes	No
Ethernet-II	Yes	Yes	Yes	Yes

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 have another service ingress policy assigned. In the default policy, no queues are defined. All traffic is mapped to the default FC, which uses one meter by default. The following table lists the characteristics of the default policy.

Table 23: Default service ingress policy ID 1 definition

Item	Definition
Meter 1	<p>1 (one) meter all unicast traffic:</p> <ul style="list-style-type: none"> Forward class: best-effort (be) mode = trtcm1 CIR, PIR = Closest (adaptation-rule) CIR = 0 PIR = max (4000000 kbps in case of a LAG with four member ports) MBS, CBS = default (values derived from applicable policy) color-mode = color-blind
Default FC (be)	<p>1 (one) flow defined for all traffic:</p> <ul style="list-style-type: none"> All traffic mapped to best-effort (be)

When using CAM-based classification and policing, available ingress CAM hardware resources can be allocated as needed, for use with different QoS classification match criteria. By default, the system allocates a single meter and 2 classification entries, so that all traffic is mapped to a single FC and the FC uses a single meter. Users can modify the resource allocation based on the need to scale the number of entries or the number of associations (that is, the number of SAPs using a policy that uses a particular match criterion).

If no resources are allocated to a particular match criterion used in the policy, the association of that policy to a SAP fails. Allocation of classification entries also allocates meter resources, which are used to implement the per-FC per-traffic type policing function. See [Resource allocation for service ingress QoS policies using CAM-based classification](#) for information about resource usage and allocation to SAP ingress policies.

2.2.4.2 Table-based classification

The 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 provide an option to use table-based classification using either IP DSCP or dot1p classification to assign both an FC and profile on SAP ingress for use with color-aware meters. See [Table-based classification using dot1p and IP DSCP for assigning FC and profile on SAP ingress for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information.

2.2.4.3 Hierarchical ingress policing

Hierarchical ingress policing allows users to specify the amount of traffic admitted into the system per SAP. It also allows users to share the available bandwidth per SAP among the different FCs of the SAP. For example, users can allow packets classified as Internet data to use the entire SAP bandwidth when other FCs do not have traffic.

Hierarchical ingress policing provides an option to configure a SAP aggregate policer per SAP on SAP ingress. Users should configure the PIR and, optionally, the burst size of the aggregate policer.

The aggregate policer monitors the traffic on different FCs and determines whether the packet is forwarded to an identified profile or dropped. The final behavior of the packet is based on the operating rate of the following items:

- per FC policer
- per SAP aggregate policer

See the command description of **aggregate-meter-rate** command in the *7210 SAS-Mxp, S, Sx, T Services Guide* for information about the final color assigned of the packet.

The meter mode "trtcm2" (RFC 4115) is used only on SAP ingress. When the SAP aggregate policer is configured, the per FC policer can be only configured in "trtcm2" mode. The meter modes "srtCM" and "trtcm1" (formerly "trtcm", as per RFC 2698) are used in the absence of an aggregate meter.



Note:

Before using a per-SAP aggregate meter/policer, meter resources must be allocated using the **config system resource-profile ingress-internal-tcam sap-aggregate-meter** command. When the amount of resources allocated for a SAP aggregate meter is changed, the node must be rebooted. The amount of resources allocated for this feature determines the number of SAPs that can use the aggregate meter/policer. See the "System Resource Allocation" section of the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information.

2.2.5 Service egress QoS policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Service egress queues are implemented at the transition from the service core network to the service access network and are available when SAP-based egress queues and shaping is enabled by using the **configure system resource-profile qos no port-scheduler-mode** command on the 7210 SAS-Mxp and the **configure system global-resource-profile qos no port-scheduler-mode** command on the 7210 SAS-R6 IMM-b and 7210 SAS-R12 IMM-b. 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



Note:

- On egress of access ports, use either port-based egress queuing and shaping, or SAP-based egress queuing and shaping for SAPs configured on the port. Use the **config>system>resource-profile>qos>port-scheduler-mode** command on the 7210 SAS-Mxp or the **config>system>global-resource-profile>qos>port-scheduler-mode** command on the 7210 SAS-R6 IMM-b and 7210 SAS-R12 IMM-b to select the mode for SAPs configured on either access ports or hybrid ports. This per-node setting affects all SAPs and access ports.
- When **port-scheduler-mode** is enabled, the software uses eight egress queues per port (access port or hybrid port). See [Service egress QoS policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information. When **port-scheduler-mode** is disabled, the software allocates eight egress queues per SAP, which are configured using service egress policies.

- The 7210 SAS-R6 IMM-c and 7210 SAS-R12 IMM-c support only port-based egress queues (with eight egress queues per port) and port-based scheduling. These platforms do not support per-SAP service egress policies.

The sub-rate capabilities and per-service scheduling control are required to make multiple services per physical port possible. Without egress shaping, it is not possible to support more than one service per port with QoS differentiation among services. 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.

Service egress QoS policies define egress queues and map FC flows to queues. The system allocates 8 (eight) queues to service egress by default. 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**

See [Queue parameters](#) for information about the parameters that can be configured for a queue.

The optional service egress QoS policy elements include specifying a remark policy that defines IEEE 802.1p priority value remarking based on FC.

The user has an option to use SAP-based marking. With SAP-based marking the remark policy defined in the SAP egress policy associated with each SAP is used to mark the packets egressing out of SAP if marking is enabled. See the [Service egress policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) and the [Remark policies](#) for more information about marking behavior and the available options. The user can also enable port-based marking, in which case the remark policy defined in the access egress policy associated with the access port determines the marking values for all the SAPs defined on the port. See [Access egress QoS policies on 7210 SAS-Mxp](#) for more information.



Note:

- Use either port-based egress queuing and shaping or SAP-based egress queuing and shaping for SAPs configured on access ports or hybrid ports. Using the **configure system resource-profile qos port-scheduler-mode** command on the 7210 SAS-Mxp or the **configure system global-resource-profile qos port-scheduler-mode** command on the 7210 SAS-R6 IMM-b and 7210 SAS-R12 IMM-b, select the per-node mode for SAPs configured on the ports.
- For Layer 2 SAPs, if remarking is enabled in the SAP egress policy and port-based marking is disabled, the dot1p values configured in the SAP egress policy are used. For Layer 3 SAPs no marking is done.
- For Layer 2 SAPs, if remarking and port-based marking are enabled in the SAP egress policy, the dot1p values configured in the SAP egress policy are used. For Layer 3 SAPs, the dot1p and DSCP values configured in the access egress policy are used. The DSCP values configured in the access egress policy are used to mark the IP traffic sent out of Layer 2 SAPs.
- If remarking is disabled for the SAP egress policy and port-based marking is enabled, IP DSCP values are marked, including for the traffic egressing the Layer 2 SAPs configured on the port. To avoid this, it is recommended to use only FC-to-dot1p values when both Layer 2 and Layer 3 SAPs are configured on the same access port.

Each queue in a policy is associated with one of the FCs. Each queue can have individual queue parameters allowing individual rate shaping of the FCs mapped to the queue. The FC per service egress packet is determined at ingress. If the packet ingressed the service on the same node, the service ingress classification rules determine the FC of the packet. If the packet is received, the FC is marked in the tunnel transport encapsulation.

The FC-to-queue map is fixed, and the queue priority is determined by the queue number, with the higher queue number having the higher priority. The user can configure a queue to be a strict queue to change the scheduling behavior for that queue. See [Schedulers on 7210 SAS-Mxp](#) and [Schedulers on 7210 SAS-R6 and 7210 SAS-R12](#) for more information.



Note:

On the 7210 SAS-Mxp, only unicast traffic sent out of RVPLS SAPs uses per-SAP egress queues. BUM traffic sent out of RVPLS SAPs uses per-port egress queues.

Service egress QoS policy ID 1 is reserved as the default service for those that do not have another service egress policy explicitly assigned. The following table lists the characteristics of the default policy.

Table 24: Default service egress policy "default" definition

Characteristic	Item	Definition
Queues	Queue 1-8	1 (one) queue defined for each traffic class
Network-Control (nc)	Queue 8	<ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • queue-mgmt = default • queue-mode = weighted • weight = 1
High-1 (h1)	Queue7	<ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • queue-mgmt = default • queue-mode = weighted • weight = 1
Expedited (ef)	Queue 6	<ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • queue-mgmt = default • queue-mode = weighted • weight = 1
High-2 (h2)	Queue 5	<ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • queue-mgmt = default

Characteristic	Item	Definition
		<ul style="list-style-type: none"> queue-mode = weighted weight = 1
Low-1 (l1)	Queue 4	<ul style="list-style-type: none"> CIR = 0 PIR = max (line rate) queue-mgmt = default queue-mode = weighted weight = 1
Assured (af)	Queue 3	<ul style="list-style-type: none"> CIR = 0 PIR = max (line rate) queue-mgmt = default queue-mode = weighted weight = 1
Low-2 (l2)	Queue 2	<ul style="list-style-type: none"> CIR = 0 PIR = max (line rate) queue-mgmt = default queue-mode = weighted weight = 1
Best-Effort (be)	Queue 1	<ul style="list-style-type: none"> CIR = 0 PIR = max (line rate) queue-mgmt = default queue-mode = weighted weight = 1

2.3 Access ingress QoS policies

This section describes the access ingress QoS policies supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone).

An access ingress QoS policy defines the ingress QoS processing for packets received on the access port when the **configure port ethernet access-ingress-qos-mode** command is set to **port-mode**.

When the **access-ingress-qos-mode** command is set to **sap-mode**, no access ingress QoS policy can be attached to the port. When the **access-ingress-qos-mode** command is set to **port-mode**, access ingress policy 1 is attached by default, and this policy can be replaced with a user-defined access ingress QoS policy.

The access ingress QoS policy defines the ingress classification rule that uses dot1p and IP DSCP values from the packet header to map traffic flows to up to eight (8) FCs and assign a profile at ingress.

In addition, an option exists to use DEI with dot1p classification for color assignment. Each FC can be associated with up to two (2) meters: one meter for unicast traffic and another for multipoint traffic (that is, broadcast, multicast and unknown-unicast traffic). The user can configure meter characteristics, such as CIR and PIR rates, committed burst size (CBS) and maximum burst size (MBS), and so on.

Define the following to configure a basic access ingress QoS policy:

- a unique service access QoS policy ID
- parameters that can be configured for a meter, as described in [Meter/policer parameters](#)
- dot1p and DSCP classification policy to map dot1p and IP DSCP values to the FC
- an optional configuration to choose either IP DSCP or dot1p values, both, or the default FC values for classification
- an optional configuration to assign an access ingress profile for use with a color-aware meter using either a dot1p, DEI with dot1p, or IP DSCP

On the 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, or 7210 SAS-Sx 10/100GE, access ingress QoS policy ID 1 is reserved as the default policy for access ports for an access ingress policy that is not explicitly assigned. On the 7210 SAS-Mxp in non-shared mode, access ingress QoS policy ID 1 is reserved as the default policy for access ports for an access ingress policy that is not explicitly assigned.

On the 7210 SAS-Mxp, an access ingress QoS policy can be shared with multiple access ports. This feature is mutually exclusive with the use of non-shared access ingress QoS policies. To share an access ingress QoS policy, the **config>system>resource-profile>qos-access-port-shared-res-mode** command must be enabled.

Define the following to configure a basic shared access ingress QoS policy:

- a unique service access QoS policy ID
- the keyword to indicate a shared policy
- parameters that can be configured for a meter, as described in [Meter/policer parameters](#)
- the **classification-criteria** command with either the **table-criteria** or **cam-criteria** parameters
- with **table-criteria** configured, the following options are available for classification:
 - a dot1p and DSCP classification policy to map dot1p and IP DSCP values to the FC
 - an optional configuration to choose either IP DSCP or dot1p values, both, or the default FC values for classification
 - an optional configuration to assign an access ingress profile for use with a color-aware meter using either a dot1p, DEI with dot1p, or IP DSCP
- with **cam-criteria** configured, users have the option to define IP criteria-based classification entries to map traffic flows to the FC

In shared mode, access ingress QoS policy ID 65536 is reserved as the default policy for access ports for an access ingress policy that is not explicitly assigned.

The following tables lists the characteristics of the default access ingress policy in non-shared mode (default policy ID 1) and shared mode (default policy ID 65536).

Table 25: Default access ingress policy ID 1 definition for 7210 SAS-Mxp (non-shared access ingress policy mode), 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE

Item	Definition
Meter 1	<p>1 (one) meter all unicast traffic:</p> <ul style="list-style-type: none"> • Mode: trtcm1 • Adaptation rule: cir closest pir closest • CIR = 0 • PIR = max • Color mode: color-aware • MBS, CBS = default kbits (values derived from applicable policy)
Meter 9	<p>1 (one) meter all multicast traffic:</p> <ul style="list-style-type: none"> • Mode: trtcm1 • Adaptation rule: cir closest pir closest • CIR = 0 • PIR = max • Color mode: color-aware • MBS, CBS = default kbits (values derived from applicable policy)
Default FC (be)	<p>1 (one) flow defined for all traffic:</p> <ul style="list-style-type: none"> • All traffic mapped to best-effort (be) • Counter mode: in-out-profile-count • dot1p-classification 1 (default dot1p classification policy) • dscp-classification 1 (default IP DSCP classification policy) • table-classification-criteria: both-dscp-dot1p • num-qos-classifiers = 4

Table 26: Shared default access ingress policy ID 65536 definition for 7210 SAS-Mxp

Item	Definition
Meter 1	<p>1 (one) meter all unicast traffic:</p> <ul style="list-style-type: none"> • Mode: trtcm1 • Adaptation rule: cir closest pir closest • CIR = 0 • PIR = max • Color mode: color-aware

Item	Definition
	<ul style="list-style-type: none"> MBS, CBS = default kbits (values derived from applicable policy)
Default FC (be)	<p>1 (one) flow defined for all traffic:</p> <ul style="list-style-type: none"> All traffic mapped to best-effort (be) Counter mode: in-out-profile-count dot1p-classification 1 (default dot1p classification policy) dscp-classification 1 (default IP DSCP classification policy) table-classification-criteria: both-dscp-dot1p num-qos-classifiers = 2

2.4 Access egress QoS policies

This section describes access egress QoS policies.

2.4.1 Access egress QoS policies on 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE

An access egress policy defines the access port queues, the FC- to-queue mapping, and the marking characteristics for the traffic egressing toward the customer on the access ports. By configuring queue shaping rates, the individual FC traffic can be managed so that each FC traffic is within SLA limits and does not impact the serviceability of other FCs.

The number of queues available per access port on different 7210 SAS platforms is the following:

- On the 7210 SAS-T, there are 8 (eight) queues always available per access port, and all FC traffic is mapped into these separate 8 (eight) queue as per the FC-to-queue-ID map. See [Table 31: Forwarding class to queue-ID map](#) for more information.
- On the 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, 16 queues are available per access port, with 8 (eight) queues allocated for unicast traffic and 8 (eight) allocated for multicast traffic. For each of the 8 (eight) FCs, unicast and multicast traffic are mapped to two different queues, for a total of 16 queues across 8 (eight) FCs. See [Table 32: Forwarding class-to-queue ID map for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#) for more information.

To define a basic access egress QoS policy, the following are required:

- unique service access QoS policy ID
- QoS policy scope of **template** or **exclusive**
- see [Queue parameters](#) for information about the parameters configured for a queue
- on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, access egress policy allows users to specify a single rate for both multicast and unicast queues, meaning that rates cannot be specified individually for the unicast and multicast queue of an FC. Instead, the rate specified for a queue (for example, queue 8) is distributed equally to the unicast queue and multicast queue associated with the FC. See [Schedulers on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#) for more information.

- remarking (for example, IEEE 802.1p value) based on FC

On the 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE, remarking of dot1p or DSCP or both bits by default is disabled and can be enabled by the **remarking** command, with options to remark dot1p, dscp, or both present under access-egress context.

The following options are available on different platforms:

- In 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE, network mode, the user is provided with an option to remark dot1p or DSCP or both.
- In 7210 SAS-T access-uplink mode, the user is provided with an option to remark both dot1p bits and IP DSCP bits.

The FC determination per service egress packet is determined at ingress. If the packet ingressed the service on the same router, the service ingress classification rules determine the FC of the packet.

The FC is determined as follows for network mode and access-uplink mode:

- In network mode, if the packet is received over a service transport tunnel on a network port, the FC is typically determined by the MPLS LSP EXP bits.
- For 7210 SAS-T in access-uplink mode, if the packet was received on a access-uplink port, the FC is determined by the dot1p bits in the outer tag of the QinQ encapsulation.

Access egress QoS policy ID 1 is reserved as the default for access ports that do not have another access egress policy explicitly assigned. The following table lists the characteristics of the default policy.

Table 27: Default access egress policy ID 1 definition for 7210 SAS-T

Characteristic	Item	Definition
Queues	Queue 1-8	1 (one) queue defined for each traffic class
Network-Control (nc)	Queue 8	• CIR=0
		• PIR=max (line rate)
		• CBS=default (values derived for optimal buffer usage)
High-1 (h1)	Queue7	• CIR=0
		• PIR=max (line rate)
		• CBS=default (values derived for optimal buffer usage)
Expedited (ef)	Queue 6	• CIR = 0
		• PIR = max (line rate)
		• CBS = default (values derived for optimal buffer usage)
High-2 (h2)	Queue 5	• CIR = 0

Characteristic	Item	Definition	
		• PIR = max (line rate)	
		• CBS = default (values derived for optimal buffer usage)	
Low-1 (l1)	Queue 4	• CIR = 0	
		• PIR = max (line rate)	
		• CBS = default (values derived for optimal buffer usage)	
Assured (af)	Queue 3	• CIR = 0	
		• PIR = max (line rate)	
		• CBS = default (values derived for optimal buffer usage)	
Low-2 (l2)	Queue 2	• CIR = 0	
		• PIR = max (line rate)	
		• CBS = default (values derived for optimal buffer usage)	
Best-Effort (be)	Queue 1	• CIR = 0	
		• PIR = max (line rate)	
		• CBS = default (values derived for optimal buffer usage)	
Flows	Default Action	All FCs are mapped to corresponding Queues and dot1p values are marked as follows:	
		In-profile	Out-profile
Network-Control (nc)		7	7
High-1 (h1)		6	6
Expedited (ef)		5	5
High-2 (h2)		4	4
Low-1 (l1)		3	3
Assured (af)		2	2

Characteristic	Item	Definition	
Low-2 (l2)		1	1

2.4.2 Access egress QoS policies on 7210 SAS-Mxp

On 7210 SAS-Mxp, the users have an option to use either port-based egress queuing and shaping or SAP-based egress queuing and shaping for SAPs configured on access ports or hybrid ports. Use the **configure system resource-profile qos port-scheduler-mode** command to select the mode for SAPs configured on all the ports of the node.

On the 7210 SAS-Mxp platforms, an access egress policy provides different functionality based on the queuing mode in use, as described in the following sections.

2.4.2.1 Access egress QoS policy for SAP-based queuing mode on 7210 SAS-Mxp

When SAP-based egress queues are in use, 7210 SAS-Mxp supports SAP-based marking for access SAPs and port-based egress marking on access ports. SAP-based marking is only supported for Layer 2 SAPs (configured in Epipe and VPLS services).

If the user enables remarking in the SAP egress policy attached to the SAP, the remark policy configured is used to mark the packets sent out of the SAP. If remarking is disabled in the SAP egress policy attached to the SAP, the remark policy configured under the access egress policy associated with the egress access port is used to mark all packets sent out of the Layer 2 SAP configured on the access port. This is known as port-based marking. Port-based marking is supported primarily for Layer 3 SAPs (configured in VPRN and IES services). SAP-based marking is not supported for Layer 3 SAPs.

On the 7210 SAS-Mxp, no explicit CLI command is provided to choose between port-based marking and SAP-based marking for Layer 2 SAPs. Choose SAP-based marking by enabling remarking in the SAP egress policy attached to the Layer 2 SAP or choose port-based marking by disabling remarking in the SAP egress policy attached to the SAP and enabling remarking in the access egress policy associated with the access port on which the Layer 2 SAP is configured.

See [Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information about marking behavior with different remark policy types.

Access egress QoS policy ID 1 is reserved as the default access egress policy. The default policy cannot be deleted or changed. The default access egress policy is applied to all access ports which do not have another access egress policy explicitly assigned. By default sap-qos-marking is enabled.

2.4.2.2 Access egress QoS policy for port-based queuing mode on 7210 SAS-Mxp

On 7210 SAS-Mxp, when port-based queues are enabled, in addition to marking values, the access egress QoS policy provides an option to define port-based queues and scheduling.

When **port-scheduler-mode** is enabled, the software uses 8 (eight) egress queues per access port or hybrid port and all the SAPs configured on the port share the 8 (eight) egress queues for traffic sent out of that port. When **port-scheduler-mode** is enabled, the queue parameters for the access port egress queues are defined using the access egress policies.

Individual queue parameters for a specific queue can be modified using the queue override feature. See [Queue overrides for access egress QoS policies on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information.

Additionally, the marking values used to mark traffic from different FCs are defined by the remark policy in the access egress policy. Per-SAP marking cannot be used when port-based queuing mode is used. See [Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information about marking behavior with different remark policy types.

On 7210 SAS-Mxp, access egress QoS policies define egress queues and map FC flows to queues, if **port-scheduler-mode** is enabled. In **port-scheduler-mode**, the system allocates 8 (eight) queues to access port egress by default. To define a basic access egress QoS policy, the following are required:

- a unique access egress QoS policy ID
- a QoS policy scope of template or exclusive

See [Queue parameters](#) for information about the parameters that can be configured for a queue.

Optional service egress QoS policy elements include specifying a remark policy that defines IEEE 802.1p priority value remarking based on the FC.

On 7210 SAS-Mxp, when port-based queuing is used, the FC-to-queue map is fixed and the queue priority is determined by the queue number, with the higher queue number having the higher priority. The user can configure a queue to be a strict queue to change the scheduling behavior for that queue. See [Schedulers on 7210 SAS-Mxp](#) for more information about scheduling.

2.4.3 Access egress QoS policies on 7210 SAS-R6 and 7210 SAS-R12

On the 7210 SAS-R6 and 7210 SAS-R12, an access egress policy allows users to define the marking values for the traffic sent out of the access ports toward the customer. Access egress QoS policies map FC flows to marking values. In addition, based on the queuing mode used on access egress, it also defines the per-port queue parameters.

2.4.3.1 Access egress QoS policies for SAP-based queuing mode



Note:

The SAP-based queuing mode is supported only on the 7210 SAS-R6 IMM-b and 7210 SAS-R12 IMM-b. It is not supported on the 7210 SAS-R6 IMM-c and 7210 SAS-R12 IMM-c.

The 7210 SAS-R6 and 7210 SAS-R12 support SAP-based marking for access SAPs and port-based egress marking on access ports. SAP-based marking is only supported for L2 SAPs, which are SAPs configured in Epipe and VPLS services. If users enable remarking in the SAP-egress policy attached to the SAP, the configured remark policy is used to mark the packets sent out of the SAP. If remarking is disabled in the SAP-egress policy attached to the SAP, the remark policy configured under the access egress policy associated with the egress-access port is used to mark all packets sent out of the L2 SAP configured on the access port. This is known as port-based marking.

Port-based marking is supported primarily for L3 SAPs (that is, SAPs configured in VPRN services and IES services). SAP-based marking is not supported for L3 SAPs.

No explicit CLI command is provided to choose between port-based marking and SAP-based marking for L2 SAPs. The user can choose SAP-based marking by enabling remarking in the SAP-egress policy attached to the L2 SAP, or port-based marking by disabling remarking in the SAP egress policy attached to

the SAP and enabling remarking in the access egress policy associated with the access port where the L2 SAP is configured.

See [Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information about marking behavior with different remark policy types.

2.4.3.2 Access egress QoS policy for port-based queuing mode



Note:

The port-based queuing mode is supported on IMM-b and IMM-c cards. On IMM-c cards, it is the default and the only supported queuing mode.

When port-based queues are enabled in addition to marking values, the access egress QoS policy provides an option to define port-based queues and scheduling.

Users have an option to use either port-based egress queuing and shaping or SAP-based egress queuing and shaping for SAPs configured on access ports or hybrid ports. The **config system global-resource-profile qos port-scheduler-mode** command allows users to select the mode used for SAPs configured on all the ports of the node (this is a per-node setting). When **port-scheduler-mode** is enabled, the software uses eight egress queues per access port, and all the SAPs configured on the port share the eight egress queues for traffic sent out of that port.

When **port-scheduler-mode** is enabled, the queue parameters for the access port egress queues are defined using the access egress policies.

Modify queue parameters for a specific queue using the queue override feature. See [QoS overrides for meters/policers](#) for more information.

Additionally, the marking values used to mark traffic from different FCs are defined by the remark policy in the access egress policy. Per-SAP marking cannot be used when port-based queuing mode is used. See [Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information about marking behavior with different remark policy types.

Access egress QoS policies define egress queues and map FC flows to queues, if **port-scheduler-mode** is enabled. Using **port-scheduler-mode**, the system allocates eight queues to access port egress by default. To define a basic access egress QoS policy, the following are required:

- unique access egress QoS policy ID
- QoS policy scope of **template** or **exclusive**

See [Queue parameters](#) for information about the parameters that can be configured for a queue.

Optional service egress QoS policy elements include the following:

- specifying a remark policy that defines IEEE 802.1p priority value remarking based on FC

When port-based queuing is used, the FC-to-queue map is fixed and the queue priority is determined by the queue number, with higher queue numbers having the higher priority. The user can configure a queue to be a strict queue to change the scheduling behavior for the queue.

Access egress QoS policy ID 1 is reserved as the default access egress policy. The default policy cannot be deleted or changed. The default access egress policy is applied to all access ports that do not have another access egress policy explicitly assigned. By default, **sap-qos-marking** is enabled. [Table 28: Default access egress QoS policy "default" definition for 7210 SAS-R6 and 7210 SAS-R12](#) and [Table 29: Default remarking policy for dot1p on 7210 SAS-R6 and 7210 SAS-R12](#) list the characteristics of the default policy.

Table 28: Default access egress QoS policy "default" definition for 7210 SAS-R6 and 7210 SAS-R12

Characteristic	Item	Definition
Network-Control (nc)	Queue 8	<ul style="list-style-type: none"> • CIR = 0 • PIR = max • queue-mgmt = default • queue-mode = weighted • weight = 1
High-1 (h1)	Queue7	<ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • queue-mgmt = default • queue-mode = weighted • weight = 1
Expedited (ef)	Queue 6	<ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • queue-mgmt = default • queue-mode = weighted • weight = 1
High-2 (h2)	Queue 5	<ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • queue-mgmt = default • queue-mode = weighted • weight = 1
Low-1 (l1)	Queue 4	<ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • queue-mgmt = default • queue-mode = weighted • weight = 1
Assured (af)	Queue 3	<ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • queue-mgmt = default • queue-mode = weighted • weight = 1

Characteristic	Item	Definition
Low-2 (l2)	Queue 2	<ul style="list-style-type: none"> CIR = 0 PIR = max (line rate) queue-mgmt = default queue-mode = weighted weight = 1
Best-Effort (be)	Queue 1	<ul style="list-style-type: none"> CIR = 0 PIR = max (line rate) queue-mgmt = default queue-mode = weighted weight = 1

The marking values used to mark traffic from different FCs are defined by the remark policy in the access egress policy. The following table lists the remarking policy for type dot1p for the 7210 SAS-R6 and 7210 SAS-R12.

Table 29: Default remarking policy for dot1p on 7210 SAS-R6 and 7210 SAS-R12

FC	In-profile	Out-profile
Network-Control (nc)	7	7
High-1 (h1)	6	6
Expedited (ef)	5	5
High-2 (h2)	4	4
Low-1 (l1)	3	2
Assured (af)	3	2
Low-2 (l2)	1	1
Best-Effort (be)	0	0

2.4.4 Queue overrides for access egress QoS policies on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

The queue override feature for access egress QoS policies allows users to override the queue parameter settings of an access egress QoS policy applied to a port. An access egress QoS policy defines the QoS behavior on access port egress. Queue override is used when port-based queues with shaping and scheduling are configured for use instead of per-SAP egress queues.

Queue override support on access egress ports allows the user to override queue parameters, such as adaptation rule, percent CIR and PIR rates, queue management policy, queue mode, CIR and PIR rates, and queue weight. See [Queue parameters](#) for parameter descriptions.

When the queue override feature is not used, queue parameters for the port are taken from the access egress QoS policy assigned to the port.

Queue override commands are supported on all Ethernet access ports.

2.4.4.1 Configuring access egress QoS policy queue override parameters

See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for CLI command descriptions of the queue override parameters.

Output example

The following is a sample queue override parameter configuration output.

```
*7210SAS>config>port>ethernet>access>egress# info
-----
      qos 13
      queue-override
        queue 4
          adaptation-rule cir closest pir closest
          queue-mgmt default
          queue-mode weighted
          rate cir 300 pir 400
          weight 5
        exit
      exit
-----
*A:7210SAS>config>service>epipe>sap>ingress#
```

2.5 Remark policies on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (network mode)

This policy allows the user to define the FC-to-egress marking values. Based on the packet encapsulation used to send out the service packets, the remark policy allows the user to define and associate policies to service egress and network egress QoS policies.

The 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (network mode) supports the following types of remark policies:

- **dot1p**
This type is used for service egress, access port egress, and network QoS (**port** type).
- **dscp**
This type is used for access port egress and network QoS (**port** type) policies.
- **lsp-exp**
This type is used for network QoS (**ip-interface** type) policies.
- **dot1p-dscp**

This type is used for access port egress and network QoS (**port** type) policies.

- **dot1p-lsp-exp-shared**

This type is used for access port egress, and network QoS (**ip-interface** type) policies.

Each of these remark policy types can be associated with only specific QoS policies, as in the preceding list.



Note:

On the 7210 SAS-R6 and 7210 SAS-R12, when the port-based scheduling mode is enabled per-node for SAPs, only per-port egress marking policies are supported; per-SAP egress marking is not supported for L2 and L3 SAPs.

The required elements to define a remark QoS policy are:

- unique remark QoS policy ID
- FC to appropriate marking values

See [Remark policies](#) for more information.

2.5.1 Egress port rate limiting

The 7210 SAS supports port egress rate limiting, which allows the user to limit the bandwidth available on the egress of the port to a value less than the maximum possible link bandwidth. It also allows the user to control the amount of burst sent out.

2.6 Forwarding classes

7210 SAS devices support multiple FCs and class-based queuing, so the concept of FCs is common to all QoS policies.

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

Queues are created for a specific FC to determine the manner in which the queue output is scheduled. The FC of the packet, along with the in-profile or out-of-profile state, determines how the packet is queued and handled (the per-hop behavior (PHB)) at each hop along its path to a destination egress point. 7210 SAS devices support 8 (eight) FCs. The following table describes the default definitions for the FCs.

Table 30: Forwarding classes

FC-ID	FC name	FC designation	DiffServ name	Notes
7	Network Control	NC	NC2	Intended for network control traffic.
6	High-1	H1	NC1	Intended for a second network control class or delay/jitter sensitive traffic.

FC-ID	FC name	FC designation	DiffServ name	Notes
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	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	Intended for BE traffic.
0	Best Effort	BE	BE	

The FC behavior, in terms of ingress marking interpretation and egress marking, can be changed by using [Network QoS policies in network mode](#). All FC queues support the concept of in-profile and out-of-profile.

2.6.1 Forwarding class-to-queue ID map on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T

There are 8 (eight) FCs supported on 7210 SAS devices. Each FC is mapped to a specific queue. By mapping FCs to different queues the differential treatment is imparted to various classes of traffic.

On the 7210 SAS devices, there are only 8 (eight) queues available at the port level for all ports. These 8 (eight) queues are created by default per port. Users cannot create or delete the queues nor the queue ID. Only the queue parameters can be changed. The queue ID is not configurable, and queue IDs 1 to 8 are, by default, used to identify the 8 (eight) queues available on the port. Queue parameters for the 8 (eight) queues can be configured using the different QoS policies based on the capabilities available on different 7210 SAS platforms. See [QoS policies](#) for more information.

The queue IDs 1 to 8 are assigned to each of the 8 (eight) queues. Queue-ID 8 is the highest priority and queue-id 1 is the lowest priority. FCs are correspondingly mapped to these queue IDs according to their priority. The following table describes the system-defined map.

Table 31: Forwarding class to queue-ID map

FC-ID	FC name	FC designation	Unicast queue-ID
7	Network control	NC	8
6	High-1	H1	7
5	Expedited	EF	6
4	High-2	H2	5
3	Low-1	L1	4

FC-ID	FC name	FC designation	Unicast queue-ID
2	Assured	AF	3
1	Low-2	L2	2
0	Best-Effort	BE	1

2.6.2 Forwarding class-to-queue ID map on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE

There are 8 FCs supported on the device. Each of these FC is mapped to two queues, one used for unicast traffic mapped to the FC and another used for multicast/BUM traffic mapped to the FC. By mapping an FC to different queues, the differential treatment is imparted to various classes of traffic.

In the 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE devices, there are 16 queues available at the port level for all ports on the device. These 16 queues are created by default per port. Users cannot create or delete the queues or the queue ID. Only the queue parameters can be changed. On 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, the queue ID represents one of the 8 scheduling nodes corresponding to 8 FCs. It is not a configurable entity and queue ID 1 to 8 are, by default, used to identify these 8 FC scheduling nodes available on the port. Parameters for these 8 FCs can be configured as part of the access egress QoS policy that is applied on the access ports, and network queue policy that is applied on the network ports and hybrid ports.

The queue IDs 1 to 8 are assigned to each of the 8 FCs. Queue ID 8 is the highest priority scheduling node and queue ID 1 is the lowest priority scheduling node. FCs are correspondingly mapped to these queue IDs according to their priority. The following table describes the system-defined map.

Table 32: Forwarding class-to-queue ID map for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE

FC-ID	FC name	FC designation	Unicast queue-ID	Multicast queue-ID
7	Network control	NC	8	8
6	High-1	H1	7	7
5	Expedited	EF	6	6
4	High-2	H2	5	5
3	Low-1	L1	4	4
2	Assured	AF	3	3
1	Low-2	L2	2	2
0	Best-Effort	BE	1	1

2.7 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 access egress QoS policy, two default network QoS policies (one each for the network QoS policy of the **ip-interface** type and of the **port** type) and two default port scheduler policies. Only one ingress QoS policy and one egress QoS policy can be applied to a SAP, access port, IP interface, network port, hybrid port, or access uplink port (the support for QoS policy association is different on different 7210 SAS platforms).

When creating a QoS policy, default values are provided for most parameters, with the exception of the policy ID and descriptions. Each policy has a scope, default action, description, meters for ingress policies, and queues for egress policies. The queue is associated with an FC.

QoS policies can be applied to the following service types on 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (based on the service supported in the particular mode of operation):

- **Epipe**
Only SAP ingress policies are supported on an Epipe service access point (SAP).
- **VPLS**
Only SAP ingress policies are supported on a VPLS SAP.
- **VPRN**
Only SAP ingress policies are supported on a VPRN SAP.
- **R-VPLS**
Only SAP ingress policies are supported on a R-VPLS SAP.

QoS policies can be applied to the following entities on the 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T:

- access egress policies and port scheduler policies on access ports
- network QoS policies on network ports and hybrid ports in network mode and on access uplink ports in access uplink mode
- network queue policies (egress) on network ports and hybrid ports in network mode and on access uplink ports in access uplink mode

On the 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T, all SAPs across all services on a specific port share the egress port queues, which can be configured using the access egress policy and port scheduler policy on an access port.

QoS policies can be applied to the following service types on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 in SAP-based egress queuing mode. SAP-egress policies are not supported when port-based egress queuing mode is configured:

- **Epipe**
SAP ingress policies and SAP egress policies are supported on an Epipe SAP.
- **VPLS**
SAP ingress policies and SAP egress policies are supported on a VPLS SAP.
- **VPRN**
SAP ingress policies and SAP egress policies are supported on a VPRN SAP.

- **IES**

SAP-ingress and SAP-egress policies are supported on an IES service.

- **R-VPLS**

SAP ingress policies and SAP egress policies are supported on a R-VPLS SAP.



Note:

SAP egress policies are not supported when port-based egress queueing mode is configured. Instead, access egress policies must be used when port-based scheduling is configured and all the SAPs on the port share the egress port queues.

QoS policies can be applied to the following entities on the 7210 SAS-R6 and 7210 SAS-R12, and 7210 SAS-Mxp:

- option to use access egress policies for marking in SAP-based scheduling mode and to configure queues and marking in port-based scheduling mode; access egress policies on access ports
- network QoS policies on network ports and hybrid ports in network mode and on access uplink ports in access uplink mode
- network queue policies (egress) on network ports and hybrid ports in network mode

QoS policies allow operators to prioritize traffic according to the FC and use congestion management to control access ingress, access egress, and network traffic (network port or access-uplink port), enqueueing packets according to their priority (color).

2.7.1 QoS policies for hybrid ports on 7210 SAS-T

This section provides an overview of QoS policy support on hybrid ports:

- Network queue policies are supported for queue configuration of egress queues on hybrid ports. These egress queues are shared by traffic sent out of SAPs and network IP interfaces configured on hybrid ports.
- Network QoS (type = **ip-interface**) policies are supported for network IP interfaces on hybrid ports. The behavior is similar to the existing behavior for network IP interfaces on network ports. It supports per IP interface ingress classification and policing, and egress marking (only EXP marking for MPLS traffic).
- Network QoS (type = **port**) policies are supported for hybrid ports. The behavior is similar to existing behavior for network ports. It supports per port ingress classification and policing, and egress marking (dot1p and/or DSCP marking) for IP control packets.
- SAP ingress QoS policies are supported for SAPs configured on hybrid ports. The behavior is similar to existing behavior for access SAP ingress. It supports per SAP ingress classification and policing.
- For marking traffic sent out of SAPs and IP traffic sent out of IP interfaces configured on hybrid ports, the user must use the network QoS policy of the **port** type, with an option to mark dot1p, DSCP, or both.



Note:

See [Remark policies](#) for more information about dot1p and IP DSCP marking.

2.7.2 QoS policies for hybrid ports on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE

This section provides an overview of QoS policy support on hybrid ports:

- Network queue policies are supported for queue configuration of egress queues on hybrid ports. The egress queues are shared by traffic sent out of SAPs and network IP interfaces configured on hybrid ports.
- Network QoS (type = **ip-interface**) policies are supported for network IP interfaces on hybrid ports. The behavior is similar to the existing behavior for network IP interfaces on network ports. It supports per IP interface ingress classification and policing, and egress marking (only EXP marking for MPLS traffic).
- Network QoS (type = **port**) policies are supported for hybrid ports. The behavior is similar to existing behavior for network ports. It supports per port ingress classification and policing, and egress marking (dot1p and/or DSCP marking) for IP control packets.
- SAP ingress QoS policies are supported for SAPs configured on hybrid ports. The behavior is similar to existing behavior for access SAP ingress. It supports per SAP ingress classification and policing.
- For marking traffic sent out of SAPs and IP traffic sent out of IP interfaces configured on hybrid ports, users must use the network QoS policy of the **port** type, with an option to mark dot1p, DSCP, or both.



Note:

See [Remark policies](#) for more information about dot1p and IP DSCP marking.

2.7.3 QoS policies for hybrid ports on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12



Note:

SAP-based egress QoS policies are not supported on the 7210 SAS-R6 IMM-c and 7210 SAS-R12 IMM-c.

This section provides an overview of QoS policy support on hybrid ports.

The following policies are available when SAP-based queues are enabled:

- Network queue policies are supported for queue configuration of egress queues on hybrid ports. These egress queues are used by traffic sent out of network IP interfaces configured on hybrid port.
- Network QoS (type = **ip-interface**) policies are supported for network IP interfaces on hybrid ports. The behavior is similar to the existing behavior for network IP interfaces on network ports. It supports per IP interface ingress classification and policing, and egress marking (only EXP marking for MPLS traffic).
- Network QoS (type = **port**) policies are supported for hybrid ports. The behavior is similar to existing behavior for network ports. It supports per port ingress classification and policing, and egress marking (dot1p or DSCP marking). For marking IP control traffic sent out of network IP interfaces configured on hybrid port, user needs to use the network QoS policy of the **port** type, with an option to mark dot1p, DSCP, or both.
- SAP ingress QoS policies are supported for SAPs configured on hybrid ports. The behavior is similar to the behavior for access SAP ingress, which supports per SAP ingress classification and policing.
- Service egress policies are supported for queue configuration of per SAP egress queues for SAPs configured on hybrid ports. These egress queues are used by traffic sent out of SAPs configured on hybrid ports. For traffic sent out of SAPs configured in Layer 2 services, dot1p can be marked per SAP

using the service egress policy. An option is provided to mark only dot1p, only IP DSCP, or both dot1p and IP DSCP, for traffic sent out of SAPs configured in Layer 3 services. One of the options can be used per port by configuring the network QoS **port** type policy.



Note:

See [Remark policies](#) for more information about dot1p and IP DSCP marking.

When port-based queues are enabled for SAPs, the QoS policies available for use with hybrid ports is the same as the preceding list, with the following exceptions:

- Network queue policies are supported for queue configuration of egress queues on hybrid ports. The egress queues are shared by traffic sent out of SAPs and network IP interfaces configured on hybrid ports, which means that per SAP service egress policies are not available for use.
- For traffic sent out of SAPs (both for SAPs configured in Layer 2 services and in Layer 3 services) configured on hybrid ports, an option is provided to mark only dot1p, only IP DSCP, or both dot1p and IP DSCP. One of the options can be used per port, by configuring the network QoS **port** type policy.

2.8 Meters/policers

This section provides information about meters/policers

2.8.1 Meter/policer parameters

This section describes the meter parameters available. Meters are available for use in both network mode and access-uplink mode with the following policies.

In network mode of operation meter/policer is available with the following:

- SAP ingress policies
- network QoS policies of the **port** type, associated with network port ingress or hybrid port ingress
- network QoS policy of the **ip-interface** type, associated with network IP interfaces configured on the network port of hybrid ports

In access-uplink mode of operation meter/policer is available with the following:

- SAP ingress policies
- network QoS policy associated with access-uplink port ingress

2.8.1.1 Meter ID

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

2.8.1.2 Committed information rate for meters

The committed information rate (CIR) for a meter is the long term average rate at which traffic is considered as conforming traffic or in-profile traffic. The higher the rate, the greater the expected

throughput. The user is able to burst above the CIR and up to PIR for brief periods of time. The time and profile of the packet is decided based on the burst sizes, as described in the following sections.

When defining the CIR for a meter, the value specified is the administrative CIR for the meter. The 7210 SAS devices have a number of native rates in hardware that are used to determine the operational CIR for the meter. 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 and PIR combination specified. See the interpretation of the administrative CIR in [Adaptation rule for meters](#).

The CIR for meters is provisioned on service ingress and network ingress within service ingress QoS policies and network QoS policies, respectively.

2.8.1.3 Peak information rate for meters

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

When defining the PIR for a meter, the value specified is the administrative PIR for the meter. The 7210 SAS devices have a number of native rates in hardware that are used to determine the operational PIR for the meter. The user has some control over how the administrative PIR is converted to an operational PIR if the hardware does not support the exact CIR and PIR combination specified. See the interpretation of the administrative PIR in [Adaptation rule for meters](#).

The PIR for meters is provisioned on service ingress and access uplink ports or network port ingress within service ingress QoS policies and network QoS policies, respectively.

2.8.1.4 Color-aware and color-blind policers

The 7210 SAS devices support color-aware policing at network ingress, whereas at service ingress a policing option is provided to use either color-aware policing or color-blind policing. In color-aware policing, users can define the color of the packet using the classification and send the colored packets to the meter. A color-aware meter treats the packets according to the color defined:

- If the packet is pre-colored as in-profile (or green packets), depending on the burst size of the packet, the meter can either mark it as **in-profile** or **out-profile**.
- If the packet is pre-colored as **out-profile** (or yellow packets), even if the packet burst is less than the current available CBS, the packet is not marked as **in-profile**, but remains as **out-profile**.
- If the packet burst is higher than the MBS, the packet is marked as red and is dropped by the meter at ingress.

The profile marked by the meter is used to determine the eligibility of the packet to be enqueued into a buffer at egress (that is, when a slope policy is configured at the egress).

In **color-blind** mode, the profile (color) assigned to the packet on ingress is ignored. The CIR and PIR rates configured for the meter determine the final profile (color) for the packet. If the packet is within the CIR, the final profile (color) assigned to the packet is in-profile (green). If the packet exceeds the CIR and is within the PIR, the final profile (color) assigned to the packet is out-of-profile (yellow). Packets that exceed the PIR rate (red) are dropped.

In **color-aware** mode, the meter uses the profile assigned to the packet on ingress. The profile can be assigned on ingress either by enabling DEI classification as done on access ports or by assigning profile based on either dot1p or DEI as done on network ports and access-uplink ports. On the 7210 SAS-Mxp,

7210 SAS-R6 and 7210 SAS-R12, the profile can also be assigned on service (or SAP) ingress by using a DSCP classification policy to set FC and profile. See [Table-based classification using dot1p and IP DSCP for assigning FC and profile on SAP ingress for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information about table-based classification.

2.8.1.5 Adaptation rule for meters

The adaptation rule provides the QoS provisioning system with the ability to adapt the administrative rates provisioned for CIR and PIR, to derive the operational rates based on the underlying capabilities of the hardware. The administrative CIR and PIR rates are translated to actual operational rates enforced by the hardware meter. The rule provides a constraint when the exact rate is not available as a result of hardware capabilities.

The following table lists the hardware rate step-size for the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T.

Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T

Rate (kbits_sec)	Burst (kbits_burst)	Rate step size (bits)	Burst step size (bits)
0 to 4194296	0 to 16773	8000	4096
4194297 to 8388592	16774 to 33546	16000	8192
8388593 to 16777184	33547 to 67092	32000	16384
16777185 to 33554368	67093 to 134184	64000	32768
33554369 to 67108736	134185 to 268369	128000	65536
67108737 to 134217472	268370 to 536739	256000	131072
134217473 to 268434944	536739 to 1073479	512000	262144
268434945 to 536869888	1073480 to 2146959	1024000	524288

The following table lists the hardware rate step-size for the 7210 SAS-Sx 10/100GE.

Table 34: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Sx 10/100GE

Rate (kbits_sec)	Burst (kbits_burst)	Rate step size (kb/s)	Burst step size (bits)
8 to 16777208	4 to 16773	8	4096
16777209 to 33554416	16774 to 33546	16	8192
33554417 to 67108832	33547 to 67092	32	16384
67108833 to 134217664	67093 to 134184	64	32768

Rate (kbits_sec)	Burst (kbits_burst)	Rate step size (kb/s)	Burst step size (bits)
134217665 to 268435328	134185 to 268369	128	65536
268435329 to 536870656	268370 to 536739	256	131072
536870657 to 1073741312	536739 to 1073479	512	262144
1073741313 to 2147482624	1073480 to 2146959	1024	524288

The system attempts to find the best operational rate depending on the defined constraint. The supported constraints are the following:

- **minimum**
Find the next multiple of step-size that is equal to or greater than the specified rate.
- **maximum**
Find the next multiple of step-size that is equal to or less than the specified rate.
- **closest**
Find the next multiple of step-size that is closest to the specified rate.

The following table lists the rate values configured in the hardware when different PIR or CIR rates are specified in the CLI.

Table 35: Administrative rate example

Administrative rate	Operation rate (min)	Operation rate (max)	Operation rate (closest)
8	8	8	8
10	16	8	8
118085	11808	11800	11808
46375	46376	46368	46376

If the user has configured any value greater than 0 and less than 648, the operation rate configured on hardware is 648 kbps, regardless of the constraint used.



Note:

The configured burst size affects the rate step-size used by the system. The system uses the step size so that both the burst-size and rate parameter constraints are met. For example, if the rate specified is less than 4 Gbps but the configured burst size is 17 Mbits, the system uses a rate step size of 16 Kbits and a burst step size of 8192 bits (see [Table 33: Supported](#)

hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T, row 2).

If the meter is a srTCM meter, both rate and burst constraints specified for both CBS and MBS are considered together to determine the step-size to use for CIR, CBS, and MBS parameters.

If the meter is a trTCM1 meter, the CIR rate and CBS burst parameters are considered together to determine the step-size to use for CIR and CBS parameters, and the PIR rate and MBS burst parameters are considered together to determine the step-size to use for PIR and MBS parameters.

If the meter is a trTCM2 meter, the CIR rate and CBS burst parameters are considered together to determine the step-size to use for CIR and CBS parameters, and the PIR (EIR) rate and MBS (EBS) burst parameters are considered together to determine the step-size to use for PIR (EIR) and MBS (EBS) parameters.

2.8.1.6 Committed burst size (for meter/policers)

The committed burst size (CBS) parameter specifies the maximum burst size that can be transmitted by the source at the CIR while still complying with the CIR. If the transmitted burst is lower than the CBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter configured parameters.

The operational CBS set by the system is adapted from the user configured value by using the minimum constraint.



Note:

See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about the **burst** parameter step-size.

2.8.1.7 Maximum burst size (for meter/policers)

For trTCM, the maximum burst size parameter specifies the maximum burst size that can be transmitted by the source at the PIR while complying with the PIR. If the transmitted burst is lower than the MBS value, the packets are marked as out-profile by the meter to indicate that the traffic is not complying with CIR, but complying with PIR.

For srTCM, the maximum burst size parameter specifies the maximum burst size that can be transmitted by the source while not complying with the CIR. If the transmitted burst is lower than the MBS value, the packets are marked as out-profile by the meter to indicate that the traffic is not complying with CIR. If the packet burst is higher than MBS, packets that are marked as red are dropped.

The operational MBS set by the system is adapted from the user-configured value by using the minimum constraint.



Note:

See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about the burst parameter step-size.

2.8.1.8 Meter counters

The 7210 SAS devices maintain counters for meters within the system for granular billing and accounting. Each meter maintains the following counters:

- counters for packets or octets marked as in-profile by the meter
- counters for packets or octets marked as out-of-profile by the meter
- optional counter for dropped and forwarded packets and octets is supported

2.8.1.9 Meter modes

The 7210 SAS devices support the following meter modes:

- srtcm - Single Rate Three Color Marking
- trtcm - Two Rate Three Color Marking
- trtcm1 - Two Rate Three Color Marking1 (applicable only for service ingress QoS policies)
- trtcm2 - Two Rate Three Color Marking2 (applicable only for service ingress QoS policies)

In srtcm, the CBS and MBS token buckets are replenished at single rate, that is, CIR. Whereas in the case of trtcm, CBS and MBS buckets are individually replenished at CIR and PIR rates, respectively. trtcm1 implements the policing algorithm defined in RFC 2698, and trtcm2 implements the policing algorithm defined in RFC 4115.

2.8.1.10 QoS overrides for meters/policers

Support for the QoS override feature on an access SAP allows the user to override the meter parameters, such as CBS, MBS, rate (CIR and PIR), mode, and adaptation rule (CIR and PIR) at the SAP context.

When meter parameter values are not overridden, the values are taken from the SAP ingress policy. That is, QoS override is not used.

Meter override commands are supported on all types of access SAP.

2.8.1.10.1 Configuration guidelines for QoS override

The configuration guidelines for QoS override are as follows:

- QoS override commands can be used only for the meters or policers defined in the SAP ingress policy.
- The 7210 SAS does not support SAP ingress queues.
- QoS override commands are not allowed when the attached policy is of an **exclusive** type.
- QoS override commands are not allowed on mirror destination SAPs.
- QoS override commands are not allowed when ToD is attached to the SAP.
- On 7210 SAS devices configured in access-uplink mode, QoS override commands are not supported for ingress and egress QoS policies used with access-uplink SAPs and ports.
- QoS override commands are not supported for ingress and egress QoS policies used with network IP interfaces and network ports.

- On the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, QoS override commands are not supported for SAP-egress queues configured in the SAP-egress QoS policies.

2.8.1.10.2 Configuring meter override parameters

Output example

The following is a sample meter override parameter configuration output.

```
*7210SAS>config>service>epipe>sap>ingress# info
-----
      qos 13
      meter-override
        meter 1 create
          mode trtcm2
          adaptation-rule pir max cir max
          cbs 300
          mbs 200
          rate cir 300 pir 400
        exit
      exit
-----
*A:7210SAS>config>service>epipe>sap>ingress#
```

2.9 Queue management

This section provides information about QoS queue management.

2.9.1 Queue parameters

This section describes the queue parameters available for queues. Queues are available for use in both network mode and access-uplink mode with the following policies.

In network mode of operation, queue is configured with the following QoS policies on 7210 SAS platforms:

- On 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T network mode:
 - network queue policies associated with network port or hybrid port egress
 - access egress policies associated with access port egress
- On 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12:
 - network queue policies associated with network port or hybrid port egress
 - SAP egress policies associated with SAP egress (when SAP-based egress queuing and scheduling is enabled for use)
 - access egress policies associated with access port egress (when port-based egress queuing and scheduling is enabled for use)

In access-uplink mode of operation, a queue is available with the following platforms:

- On the 7210 SAS-T access-uplink mode:
 - network queue policies associated with access-uplink port egress

- access egress policies associated with access port egress

2.9.1.1 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 where the queue is defined. On the 7210 SAS, the queue ID is not a user configurable entity, but the queue ID is statically assigned to the 8 queues on the port according to the FC-QID map listed in [Table 30: Forwarding classes](#).

2.9.1.2 CIR for queues

The CIR for a queue performs the following distinct functions:

- **minimum bandwidth guarantees**

The egress queue CIR setting provides the bandwidth for this queue as compared to other queues on the port competing for a share of the available link bandwidth. The queue CIR does not necessarily guarantee bandwidth in all scenarios and also depends on factors such as CIR over-subscription and link port bandwidth capacity. For each packet in an egress queue, the CIR is checked with the current transmission rate of the queue. If the current rate is at or below the CIR threshold, the queue is considered in-profile. If the current rate is above the threshold, the queue is considered out-of-profile. This in and out profile state of queue is linked to scheduler prioritizing behavior as discussed in the following point.

- **scheduler queue priority metric**

The scheduler that serves a group of egress queues prioritizes individual queues based on the current CIR and PIR states. Queues operating below their CIR are always served before those queues operating at or above their CIR. See [QoS port scheduler policies for 7210 SAS-T](#), [Schedulers on 7210 SAS-Mxp](#), and [Schedulers on 7210 SAS-R6 and 7210 SAS-R12](#) for information about scheduler behavior.

Queues at the egress never mark the packets as in-profile or out-profile based on the queue CIR and PIR values. The in-profile and out-profile state of the queue interacts with the scheduler mechanism and provides the minimum and maximum bandwidth guarantees.

When defining the CIR for a queue, the value specified is the administrative 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 and PIR combination specified. See [Adaptation rule for queues](#) for information about the interpretation of the administrative CIR.

Although the 7210 SAS is flexible in how the CIR can be configured, there are conventional ranges for the CIR based on the FC of a queue. A access egress queue associated with the high-priority class normally has the CIR threshold equal to the PIR rate, although the 7210 SAS allows the CIR to be provisioned to any rate below the PIR if this behavior is required.

The CIR for a queue is provisioned in the appropriate queue policy associated with the service object (that is, a network or hybrid port, or an access SAP, as applicable).

2.9.1.3 PIR for queues

The PIR defines the maximum rate at which packets are allowed to exit the queue. It does not specify the maximum rate at which packets may enter the queue; this is determined by the ability of the queue to

absorb bursts. The actual transmission rate of an egress queue depends on more than just its PIR. Each queue is competing for transmission bandwidth with other queues. Each queue PIR, CIR, and the relative priority and weight of the scheduler serving the queue, all combine to affect a queue's ability to transmit packets.

When defining the PIR for a queue, the value specified is the administrative PIR for the queue. The user has some control over how the administrative PIR is converted to an operational PIR if the hardware does not support the exact CIR and PIR values specified. See [Adaptation rule for queues](#) for information about the interpretation of the administrative PIR.

The PIR for a queue is provisioned in the appropriate queue policy associated with the service object (that is, a network, hybrid, or access port, or an access SAP, as applicable).

2.9.1.4 Adaptation rule for queues

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

For the CIR and PIR parameters individually, the system attempts 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 less than or equal to 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 actual operational CIR and PIR settings used by the queue depend on the method the hardware uses to implement and represent the mechanisms that enforce the CIR and PIR rates.

The 7210 SAS uses a single rate step value to define the granularity for both the CIR and PIR rates. The adaptation rule controls the method the system uses to choose the rate step based on the administrative rates defined by the **rate** command.

[Table 36: Supported hardware rates and CIR/PIR values for 7210 SAS-T and 7210 SAS-Sx/S 1/10GE devices](#) lists the supported hardware rate steps that correspond to the CIR and PIR ranges between 0 and 1 Gb/s on 7210 SAS-T and 7210 SAS-Sx/S 1/10GE devices. [Table 37: Supported hardware rates and CIR/PIR values for 7210 SAS-Mxp](#) lists the supported hardware rate steps that correspond to the CIR and PIR range values between 0 to 1 Gb/s for the 7210 SAS-Mxp. [Table 38: Supported hardware rates for CIR and PIR values for 7210 SAS-R6 and 7210 SAS-R12](#) lists the supported hardware rate steps that correspond to the CIR and PIR range values between 0 to 1 Gb/s for the 7210 SAS-R6 and 7210 SAS-R12. [Table 39: Supported hardware rates and CIR/PIR values for 10-Gig port for all platforms](#) lists the supported hardware rate steps that correspond to the CIR and PIR ranges between 0 to 10 Gbps on 10-Gig ports on all platforms. [Table 40: Supported hardware rates and CIR/PIR values for 7210 SAS-Sx 10/100GE](#) lists the supported hardware rate steps that correspond to the CIR and PIR range values between 0 to 1 Gb/s for the 7210 SAS-Sx 10/100GE.

Table 36: Supported hardware rates and CIR/PIR values for 7210 SAS-T and 7210 SAS-Sx/S 1/10GE devices

Hardware rate steps	Rate range
8 kbps	0 to 16770 kbps
16 kbps	16780 to 33540 kbps
32 kbps	33550 to 67090 kbps
64 kbps	67100 to 134180 kbps
128 kbps	134190 to 268360 kbps
256 kbps	268370 to 536730 kbps
512 kbps	536740 to 1000000 kbps

Table 37: Supported hardware rates and CIR/PIR values for 7210 SAS-Mxp

Hardware rate steps	Rate range
8 kbps	0 to 16383 kbps
16 kbps	16384 to 32767 kbps
32 kbps	32768 to 65535 kbps
64 kbps	65536 to 131071 kbps
128 kbps	131072 to 262143 kbps
256 kbps	262144 to 524287 kbps
512 kbps	524288 to 1000000 kbps

Table 38: Supported hardware rates for CIR and PIR values for 7210 SAS-R6 and 7210 SAS-R12

Hardware rate steps	Rate range
8 kbps	0 to 2047 kbps
8 kbps	2048 to 4095 kbps
8 kbps	4096 to 8191 kbps
8 kbps	8192 to 16383 kbps
16 kbps	16384 to 32767 kbps
32 kbps	32768 to 65535 kbps
64 kbps	65536 to 131071 kbps

Hardware rate steps	Rate range
128 kbps	131072 to 262143 kbps
256 kbps	262144 to 524287 kbps
512 kbps	524288 to 1048575 kbps
1024 kbps	1048576 to 2097151 kbps
2048 kbps	2097152 to 4194303 kbps
4096 kbps	4194304 to 8388607 kbps
8192 kbps	8388608 to max

Table 39: Supported hardware rates and CIR/PIR values for 10-Gig port for all platforms

Hardware rate steps	Rate range
8 kbps	0 to 16383 kbps
16 kbps	16384 to 32767 kbps
32 kbps	32768 to 65535 kbps
64 kbps	65536 to 131071 kbps
128 kbps	131072 to 262143 kbps
256 kbps	262144 to 524287 kbps
512 kbps	524288 to 1048575 kbps
1024 kbps	1048576 to 2097151 kbps
2048 kbps	2097152 to 4194303 kbps
4096 kbps	4194304 to 8388607 kbps
8192 kbps	8388608 to 10000000 kbps

Table 40: Supported hardware rates and CIR/PIR values for 7210 SAS-Sx 10/100GE

Hardware rate steps	Rate range
8 kbps	8 to 16777208 kbps
16 kbps	16777209 to 33554416 kbps
32 kbps	33554417 to 67108832 kbps
64 kbps	67108833 to 134217664 kbps
128 kbps	134217665 to 268435328 kbps

Hardware rate steps	Rate range
256 kbps	268435329 to 536870656 kpbs
512 kbps	536870657 to 1073741312 kpbs
1024 kbps	1073741313 to 2147482624 kpbs

To illustrate how the adaptation rule constraints of **minimum**, **maximum**, and **closest** are evaluated in determining the operational CIR or PIR for the 7210 SAS, assume there is a queue where the administrative CIR and PIR values are 90 Kbps and 150 Kbps, respectively.

If the adaptation rule is **minimum**, the operational CIR and PIR values are 96 Kbps and 152 Kbps respectively, as the native hardware rate is greater than or equal to the administrative CIR and PIR values.

If the adaptation rule is **maximum**, the operational CIR and PIR values are 88 Kbps and 144 Kbps.

If the adaptation rule is **closest**, the operational CIR and PIR values are 88 Kbps and 152 Kbps, respectively, as those are the closest matches for the administrative values that are even multiples of the 8 Kbps rate step.

2.9.1.5 CBS and MBS for queues

The CBS and MBS parameters configure the amount of buffers that a queue can use. The CBS parameter specifies the amount of buffer reserved for a queue in the queue buffer pool. The MBS parameter specifies the portion that a queue can contend for in the shared buffer space. When all reserved buffers for a specific queue are used, the queue contends with other queues for additional buffer resources up to the configured MBS.

On the 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T, the CBS parameter for queues is not configurable for access, network, and access-uplink ports. The CBS is set to system-defined values.

On 7210 SAS-Mxp, and 7210 SAS-R6 and 7210 SAS-R12 equipped with IMM-b cards, the CBS and MBS values for the queues are configurable for the service egress and network port queues. The CBS and MBS is set to default values that address the specific FC needs to maintain differential treatment.

On the 7210 SAS-T (network and access-uplink mode), the node can be operated with either a per-port or per-node MBS pool. The decommissioning feature is supported in the per-port MBS pool mode, which increases the per-port MBS pool by taking away packet buffers from other ports. In this case, the maximum MBS per queue, assuming no other queue has traffic on that port, depends on the user configuration. For example, assuming one port is decommissioned and its buffers are allocated to port 1/1/1, the maximum MBS per queue on port 1/1/1, assuming no other queues have any traffic, is 93 Kbytes. The **show pools port-id network-egress**, **show pools port-id access-egress**, and **show pools port-id access-uplink-egress** CLI commands display the values in use depending on the port mode (network, access, and access-uplink).

The following table lists the default CBS and MBS values for the 7210 SAS platforms.

Table 41: Default CBS and MBS values

Platform	CBS (KBytes)		MBS (KBytes)	
	Network queue and access-uplink queue	Access queue	Network queue and access-uplink queue	Access queue
7210 SAS-T (MBS Pool = Port) ⁹	3.375	3.375	33 ¹⁰	33 ¹⁰
7210 SAS-T (MBS Pool = Node) ^{9 11}	1.68	1.68	1458 ¹²	1458 ¹²
7210 SAS-Mxp	128	10	256	64
7210 SAS-R6 (IMM-b)	128	10	256	64
7210 SAS-R12 (IMM-b)	128	10	256	64
7210 SAS-R6 IMM-c ¹³	1.6	1.6	1255 ¹⁴	1255 ¹⁴
7210 SAS-R12 IMM-c ¹³	1.6	1.6	1150 ¹⁴	1150 ¹⁴
7210 SAS-Sx/S 1/10GE 24-port variant ¹³	1.6	1.6	905 ¹⁴	905 ¹⁴
7210 SAS-Sx/S 1/10GE 48-port variant ¹³	1.6	1.6	698.34 ¹⁴	698.34 ¹⁴
7210 SAS-Sx 10/100GE ¹³	1.6	1.6	4158.4 ¹⁴	4158.4 ¹⁴
7210 SAS-Sx/S 1/10GE (standalone-VC) ¹³	1.6	1.6	712.7 ¹⁴	712.7 ¹⁴

2.9.2 Buffer pools

The available buffer space is partitioned into buffer pools. The buffers for a queue are allocated from the available buffer pool.

This section provides information about buffer pools for 7210 SAS platforms.

⁹ Supported both in network mode and access-uplink mode

¹⁰ The maximum MBS per queue, assuming no other queues on the same port have traffic

¹¹ Supports decommissioning feature

¹² The maximum MBS per queue, assuming only some queues have traffic

¹³ One buffer pool is shared among all queues

¹⁴ The maximum MBS per queue, assuming no other queues have traffic

2.9.2.1 Buffer pools on 7210 SAS-T

The 7210 SAS devices, when operating in network mode and access-uplink mode, support either one or both of the two modes of buffer pool allocation for port egress queues - per port MBS pool and per node MBS pool. The buffer pools take care of the buffer requirements at the port level for various queue shaping/scheduling mechanisms. In addition, in per port MBS pool mode, an option is provided to decommission the port and allocate its buffers toward other ports. The following sections provides more information about these two modes.

2.9.2.1.1 Buffer pool allocation - per port MBS pool (7210 SAS-T)

When the decommission entries are not configured, during system initialization, based on the maximum number of ports supported on the device, the total buffer is distributed into per port egress buffer pool for access ports, network ports, access-uplink ports and hybrid ports. Each port on the system gets an equal portion of the available buffers. From the buffers allocated to a port, each queue gets its CBS amount of buffers. The remaining buffers are allocated toward the shared MBS pool per port. All the queues of the port can use the buffers from the shared MBS pool. This model of buffer pool allocation is called per port MBS pool.

With the per port MBS pool, each queue is allocated with a small fixed amount of buffers toward the CBS (Committed burst size) and each port is allocated with a shared pool of buffers toward the MBS (Maximum Burst Size). The queue's CBS portion of buffers guarantees that the queue does not starve because of lack of buffers.

The buffers allocated toward the MBS pool, allow each port to handle some amount of burst. Per port MBS pool/portion of buffers is shared by all the queues of the port and allows any queue or a small group of queues of the port to absorb larger bursts assuming that, not all the queues receive burst simultaneously. In a typical network, the router/switch in the ingress traffic is usually a mix of packets of different sizes and different flows burst at different time intervals, which allows for better burst absorption capability per queue using shared resources.

2.9.2.1.2 Buffer pool allocation - per node MBS pool (7210 SAS-T)

In the per node MBS pool mode, each of the queue on a port, is allocated a CBS amount of committed buffers. The remaining amount of buffers is allocated toward the MBS pool that is available for sharing among all the queues across all the ports of the node. The queue's CBS portion of buffers guarantees that the queue does not starve because of lack of buffers.

The buffers allocated toward the node's MBS pool, allows each port to handle some amount of burst. Per port MBS pool of buffers is shared by all the queues of the port and allows any queue or a group of queues across multiple ports to absorb larger bursts, assuming that not all the queues on all the ports receive burst simultaneously. In a typical network, the router/switch in the ingress traffic is usually a mix of packets of different sizes and different flows burst at different time intervals, which allows for better burst absorption capability per queue using shared resources.

The hardware implements an algorithm to handle requests for allocation of buffers from the MBS pool (in both models) assuming that not all the ports and queues burst at the same time. This allows some queues to use a larger portion of the buffers when it is available, allowing them to handle larger bursts. At the same time, the algorithm ensures that all the queues get fair share of the buffers, so that the throughput on those ports is not affected.

When hardware receives a packet, before it decides to queue up the packet on the egress queue of the destination port, it determines the discard threshold for the queue based on the oversubscription factor and the total amount of free buffers available at that point of time.

The queue's discard threshold is higher if the amount of free buffers available is larger (which indicates other queues on the node have lesser congestion), allowing the queue to absorb larger bursts. The queue's discard threshold is lower if there are fewer free buffers available (which indicates that other queues are heavily congested on the node), which results in the packet being dropped. At the same time, algorithm allocates the available free buffers to queues which are using lesser amount of buffers or not using any buffers. This allows equal sharing of available buffers and maintains a good throughput for less congested queues.

On 7210 SAS-T (in both access-uplink and network mode), 2MB of buffers are available and by default per node MBS pool is used and an option is available to user to change it to per port MBS pool.



Note:

- On the 7210 SAS-T (network mode and access-uplink mode) with either a per-port MBS pool or a per-node MBS pool, system internal ports — such as an internal loopback port used for mirroring, port loopback with mac-swap, and others — are allocated buffers. Some buffers are reserved for internal use.
- Buffer pools cannot be created or deleted on the 7210 SAS.

2.9.2.2 Decommissioning ports with per port MBS pool

To allow operators better control over which ports get larger portion of queue buffers, the operator is provided with an option to use per-port MBS pool and decommission ports. The decommissioning of ports is only allowed when the node is booted with the option to use per-port MBS pool.

With the decommissioning feature, the user is provided with an option to make efficient use of the available port egress queue buffer pool by allocating queue buffers of the unused ports to in-use ports. It allows the user to specify the unused front-panel ports which cannot be used to deploy any services. The software does not allocate any queue buffers to these unused ports and assigns it to a specific port or a group of ports. The user is provided with a CLI command to decommission a port and make it unavailable to deploy services. This mechanism allows operators who use limited number of ports to deploy services, to increase the amount of queue buffers allocated to them by decommissioning ports that will not be used to deploy any services.

2.9.2.2.1 Using decommission command for buffer allocation on 7210 SAS-T devices



Note:

- Using the **decommission** command for buffer allocation is only supported on the 7210 SAS-T.
- This feature is not supported on the 7210 SAS-Mxp.

This feature enables the user to make efficient use of the available port egress queue buffer pool by allocating queue buffers of the unused ports to ports. Services cannot be configured on the unused ports as software takes away all the queue buffer resources from these ports that need increased amount of buffers to handle larger bursts. This allows the operators who use limited number of ports to deploy services, to increase the amount of queue buffers allocated to them by decommissioning ports that are not used to deploy services.

The amount of credit of queue buffers received by a port is used to increase the MBS portion of the buffer pool of the port. This allows any queue on the port to use the buffers, if needed. The CBS portion of the queue is not modified with this feature.



Note:

The system has to be rebooted after decommissioning of ports for the queue buffers to be reallocated and the configuration to take effect.

The users have an option to specify the groups of ports which receive the credit of queue buffers freed up using the decommission command. With this option, the user can specify a port or group of ports which receives the credit of queue buffers. For example, it is possible for the user to configure decommissioning of 4 fixed copper ports and allocate the freed queue buffers to the remaining copper ports in the system or decommission 5 fiber ports and allocate the freed up queue buffers to the 10G XFP ports, and so on. This mechanism allows the operators to provide higher amount of queue buffers to a specific port or a group of ports, allowing them to pick and choose ports that need the extra buffers for burst absorption.

The user is allowed to increase the per port MBS pool limit so that more buffers are available to absorb larger bursts, at the cost of decommissioning ports which are not used to configure services.

2.9.2.2.2 Configuration guidelines for use of decommission commands on 7210 SAS-T devices

The **configure system resource-profile decommission entry** CLI command allows the user to configure the list of ports to be decommissioned and the list of ports that need more buffers. The system does not allocate any packet buffers to the ports which are decommissioned. For more information, see the CLI command description for details on the functionality provided by the command.

Packet buffers are added to the MBS pool of the port (the MBS pool is shared by the 8 (eight) queues on the port) and the CBS portion of the queues are not modified.

The user can modify the list of ports or update to the list of ports specified with the decommission command (and also entry command) when the node is up.

The software maintains two lists of entries, one is the current list of ports and another which has been modified by the user and takes effect only after the next reboot. These lists can be displayed using the show command. The configuration file always stores the list of entries as configured by the user, so that when rebooted the modified entries and new entries (if any) takes effect.

A port must be in administrative down (shutdown) state before it is in a decommission entry. An attempt to configure a port which is administratively up (no shutdown) state results in an error. The administrative state or the operational state of the port is not affected by configuring the port in a decommission entry.

The decommissioned port cannot be used in any service configuration or as a loopback port. An attempt to do so results in an error.

The decommissioned port must not be configured with BOF parameter, 'no-service-ports'.

Buffer allocation to a port should be possible for access ports, network ports or hybrid ports. In other words, irrespective of port mode, it is possible to assign more buffer resources to the port.

The user needs to ensure that enough buffers are available for the internal loopback ports or front-panel ports assigned for loopback. It is not recommended to take away buffers allocated to these ports and assign it to other ports. This may cause unintended behavior of the system. The system software does not check for this, but expects users to ensure this through proper configuration.

During system boot up, while executing the commands in the configuration file software checks if the no-service-ports are configured under the decommission entries. If there is match, software throws an error and stops execution of further commands in the configuration file. When this happens, user needs to correct the configuration file or the BOF file, to either remove the ports from the decommission entries or not configure them as no-service-ports in the BOF, save the BOF file or the configuration file based on where the change was made and reboot the node.

On the 7210 SAS-T, the decommission command takes affect only if the per port MBS pool is in use, that is, the user needs to configure the **configure system resource-profile qos mbs-pool port** CLI command, before using the decommission port feature.

Example

The following configuration sample shows the ports to be decommissioned and the ports that need more buffers.

```
A:7210SAS>config>system>res-prof>decom# info detail
-----
entry 15 port 1/2/1,1/2/2 to 1/1/2
entry 23 port 1/1/5 to 1/1/3
-----
A:7210SAS>config>system>res-prof>decom#
```



Note:

See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information about the decommission commands.

2.9.2.3 Buffer pools on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE

The 7210 SAS-Sx/S 1/10GE has 4 MB of buffers per node and the 7210 SAS-Sx 10/100GE has 16 MB of buffers per node (CBS and MBS); both platforms support only a single mode of operation per node MBS pool. In this mode, the MBS pool is shared across all queues on all ports. In the per node MBS pool mode, each of the 16 egress queues available on a port, is allocated a CBS amount of committed buffers. The remaining amount of buffers is allocated toward the per node MBS pool that is available for sharing among all the queues across all the ports of the node.



Note:

The system internal ports, such as internal loopback ports used for mirroring, port loopback with mac-swap, and others, are allocated with some buffers. Some buffers are reserved for system internal use (for example, CPU queues).

The amount of buffers remaining after allocating buffers for system internal use is available for allocation toward MBS buffers for all egress queues and per node MBS pool.

The hardware implements an algorithm to handle requests for allocation of buffers from the MBS pool assuming that not all the ports and queues burst at the same time. This allows some queues to use a larger portion of the buffers when it is available, allowing them to handle larger bursts. At the same time, the algorithm ensures that all the queues get fair share of the buffers, so that the throughput on those ports are not affected. When the hardware receives a packet, before it decides to queue up the packet on the egress queue of the destination port, it determines the discard threshold for the queue based on the oversubscription factor and the total amount of free buffers available at that point of time.

The queue's discard threshold is higher, if the amount of free buffers available is larger (which indicates other queues on the node have lesser congestion), allowing the queue to absorb larger bursts. The

queue's discard threshold is lower, if there is lesser amount of free buffers available (which indicates that other queues are heavily congested on the node), which results in the packet being dropped. At the same time, algorithm allocates the available free buffers to queues which are using lesser amount of buffers or not using any buffers. This allows equal sharing of available buffers and maintains a good throughput for less congested queues.



Note:

- The 7210 SAS-Sx/S 1/10GE does not support per-port MBS pool mode and port decommissioning features in this release.
- Buffer pools cannot be created or deleted on the 7210 SAS.

2.9.2.4 Buffer pools on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

The 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 have a single buffer pool per node, the system pool. All the queues created by the system are allocated buffers from this system pool. Queues come up with default buffers, and the buffers change accordingly when they are associated with a network port or SAP. Queue management policies allow the user to specify the parameters that determine buffer allocation to the queues. Buffer pools cannot be created or deleted in the 7210 SAS.

2.9.3 Queue management policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12



Note:

The 7210 SAS-R6 IMM-c and 7210 SAS-R12 IMM-c do not support the user configuration of CBS and MBS queue parameters. These values are system-defined and configured values in queue management policies are ignored. On the 7210 SAS-R6 IMM-c and 7210 SAS-R12 IMM-c, only the WRED slope parameters are used from the queue management policy. References to allocation of queue buffers in this section applies only to the 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, and 7210 SAS-Mxp.

Queue management policies allows the user to define the queue buffer and WRED slope parameters. The device supports a single buffer pool per node. All the queues created in the system are allocated buffers from this system pool. The default buffers are allocated to the queues accordingly when they are associated with a SAP or a network port.

Queue management policies allow the user to define the CBS, MBS and WRED parameters for use by the queue. The CBS and MBS parameters are used to allocate the appropriate amount of buffers from the system pool to the queues. The WRED parameters allow the user to define the WRED slope characteristics.

You can define a high-slope and a low-slope for each of the queues. High-slope is used for in-profile packets being enqueued into the queues and low-priority slope is used for out-of-profile packets being enqueued into the queues. By default each queue is associated with a default queue-management policy. The default policy allocates the appropriate amount of CBS and MBS buffers based on whether the queue is associated with a SAP or network port.



Note:

If WRED is not configured, taildrop is used.

2.9.3.1 Queue management policy parameters

The elements required to define a queue management policy are:

- a unique policy ID
- CBS and MBS parameters to allocate buffers to the queues
- the RED slope shapes for the buffer-pool, that is, start-average, max-average, max-drop-probability settings for the high-priority and low priority RED slope
- the TAF factor for the queue

Queue management policy ID default is reserved for the default queue management policy. The default policy cannot be deleted or changed. The default policy is implicitly applied to all queues that do not have another queue management policy explicitly assigned. The following table lists the default values for the default slope policy.

Table 42: Default values for the default slope policy for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Parameter	Description	Setting
Policy ID	Queue management policy ID	Default (for default queue management policy)
CBS	Committed Burst size	Default (in kilobytes)
MBS	Maximum Burst size	Default (in kilobytes)
High (RED) slope	Administrative state	Shutdown
	start-avg	70% utilization
	max-avg	90% utilization
	max-prob	75%
Low (RED) slope	Administrative state	Shutdown
	start-avg	50% utilization
	max-avg	75% utilization
	max-prob	75%
TAF	time-average-factor	7

See [Table 41: Default CBS and MBS values](#) for the default CBS and MBS queues on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

2.9.4 RED slopes in network and access-uplink mode

On 7210 SAS platforms RED slopes support is as follows:

- On 7210 SAS-T (both network mode and access-uplink mode), each queue, supports a high-priority RED slope and a low-priority RED slope.
- On 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, each queue supports a high-priority RED slope and a low-priority RED slope, which are configurable using the queue-management policies.
- On 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, each queue 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.

By default, all slopes are disabled.

The WRED uses average queue lengths, queue thresholds provisioned, and drop probability to calculate the packet's eligibility to be enqueued. The committed portion of the buffer pool is exclusively used by a queue to enqueue traffic within committed rate.

For the queues within a buffer pool, packets are either queued using committed burst size (CBS) buffers or shared buffers. The CBS buffers are simply buffer memory that has been allocated to the queue while the queue depth is at or below its CBS threshold.

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


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

1. The RED function keeps track of shared buffer utilization and shared buffer average utilization.
2. At initialization, the utilization is 0 (zero) and the average utilization is 0 (zero).
3. When each packet is received, the current average utilization is plotted on the slope to determine the packet's discard probability.
4. A random number is generated associated with the packet and is compared to the discard probability.
5. The lower the discard probability, the lower the chances are that the random number is within the discard range.
6. If the random number is within the range, the packet is discarded which results in no change to the utilization or average utilization of the shared buffers.
7. A packet is discarded if the utilization variable is equal to the shared buffer size or if the used CBS (actually in use by queues, not just defined by the CBS) is oversubscribed and has stolen buffers from the shared size, lowering the effective shared buffer size equal to the shared buffer utilization size.
8. If the packet is queued, a new shared buffer average utilization is calculated using the time-average-factor (TAF) for the buffer pool. The TAF describes the weighting between the previous shared buffer average utilization result and the new shared buffer utilization in determining the new shared buffer average utilization. (See [Tuning the shared buffer utilization calculation](#).)
9. The new shared buffer average utilization is used as the shared buffer average utilization next time a packet's probability is plotted on the RED slope.

10. When a packet is removed from a queue (if the buffers returned to the buffer pool are from the shared buffers), the shared buffer utilization is reduced by the amount of buffers returned. If the buffers are from the CBS portion of the queue, the returned buffers do not result in a change in the shared buffer utilization.

A RED slope itself is a graph with an X (horizontal) and Y (vertical) axis. The X-axis plots the percentage of shared buffer average utilization, going from 0 to 100 %. 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 (as shown in the following figure):

Figure 4: RED slope characteristics



OSSG020

1. Section A is (0, 0) to (**start-avg**, 0). This is the part of the slope that the packet discard value is always zero, preventing the RED function from discarding packets when the shared buffer average utilization falls between 0 and **start-avg**.
2. Section B is (**start-avg**, 0) to (**max-avg**, **max-prob**). This part of the slope describes a linear slope where packet discard probability increases from zero to **max-prob**.
3. Section C is (**max-avg**, **max-prob**) to (**max-avg**, 1). This part of the slope describes the instantaneous increase of packet discard probability from **max-prob** to one. A packet discard probability of 1 results in an automatic discard of the packet.
4. Section D is (**max-avg**, 1) to (100%, 1). On this part of the slope, the shared buffer average utilization value of **max-avg** to 100% results in a packet discard probability of 1.

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

2.9.4.1 Tuning the shared buffer utilization calculation

The 7210 SAS allows tuning the calculation of the Shared Buffer Average Utilization (SBAU) after assigning buffers for a packet entering a queue as used by the RED slopes to calculate a packet's drop probability. It implements a time average factor (TAF) parameter in the buffer policy which determines the contribution of the historical shared buffer utilization and the instantaneous Shared Buffer Utilization (SBU) in calculating the SBAU. The TAF defines a weighting exponent used to determine the portion of the shared buffer instantaneous utilization and the previous shared buffer average utilization used to calculate the new shared buffer average utilization. To derive the new shared buffer average utilization, the buffer pool takes a portion of the previous shared buffer average and adds it to the inverse portion of the

instantaneous shared buffer utilization (SBU). The formula used to calculate the average shared buffer utilization is shown in the following figure.

Figure 5: Calculation for average shared buffer utilization

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

sw0533

where:

$SBAU_n$ = Shared buffer average utilization for event n

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

SBU = The instantaneous shared buffer utilization

TAF = The time average factor

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

Table 43: TAF impact on shared buffer average utilization calculation

TAF	2^{TAF}	Equates to	Shared buffer instantaneous utilization portion	Shared buffer average utilization portion
0	2^0	1	1/1 (1)	0 (0)
1	2^1	2	1/2 (0.5)	1/2 (0.5)
2	2^2	4	1/4 (0.25)	3/4 (0.75)
3	2^3	8	1/8 (0.125)	7/8 (0.875)
4	2^4	16	1/16 (0.0625)	15/16 (0.9375)
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)

TAF	2^{TAF}	Equates to	Shared buffer instantaneous utilization portion	Shared buffer average utilization portion
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.

2.9.4.2 Slope policies for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE devices



Note:

On 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, queue management policies are used to configure the WRED slopes. See [Queue management policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information.

Slope policies define the RED slope characteristics as a percentage of the size of the queue on which the policy is applied when the per-port MBS pool configuration is used. When per node MBS pool is in use, the slope parameters are interpreted as a percentage of the logical size for the queue and is not a percentage of the total MBS pool size.

On 7210 SAS-T (network mode and access-uplink mode), default buffer pools exist (logically) at the port levels, when configured to use per port MBS pool and is dependent on the physical port mode:

- access egress pool on access ports
- network egress pool (in network mode) on network ports and hybrid ports
- access uplink egress pool (in access uplink mode) on access ports

By default, each queue on the port is associated with slope-policy default which disables the **high-slope**, **low-slope**, and non-TCP slope within the pool.

On the 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 1/10GE and 7210 SAS-Sx 10/100GE configured in standalone and standalone-VC mode, default buffer pools exist (logically) at the port levels and is dependent on the port mode:

- access egress pool on access ports

- network egress pool on network ports and hybrid ports

By default, each queue on the port is associated with slope-policy default, which disables the **high-slope** and **low-slope**.



Note:

If WRED is not configured, taildrop is used.

2.9.4.2.1 Slope policy parameters

The elements required to define a slope policy are:

- a unique policy ID
- on 7210 SAS-T, 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, only two slopes per queue is available
- the high and low RED slope shapes for the buffer pool: settings for the high-priority and low-priority RED slopes

All slopes are available per queue and the following parameters are configurable for each slope:

- start-avg
- max-avg
- max-prob
- Time average factor (TAF)

A slope policy is defined with generic parameters so that it is not inherently an access or a network policy. A slope policy defines access egress buffer management properties, when it is associated with an access port buffer pool and network egress buffer management properties, when it is associated with a network port buffer pool.

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 and network buffer pools which do not have another slope policy explicitly assigned.

The following table lists the default values for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE in network mode.

Table 44: Default slope policy definition (for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE configured in network mode)

Parameter	Description	Setting
Policy ID	policy ID	default (for default policy)
High (RED) slope	Administrative state	Shutdown
	start-avg	70% utilization
	max-avg	90% utilization
	max-prob	75%
Low (RED) slope	Administrative state	Shutdown

Parameter	Description	Setting
	start-avg	50% utilization
	max-avg	75% utilization
	max-prob	75%

2.9.5 CPU queues

The packets that are destined for the CPU are prioritized based on the application. Some of the applications that are prioritized are Layer 2 data packets (a copy of which is sent to CPU for MAC learning), EFM, CFM, STP, LACP, ICMP, and so on. The CPU provides 8 (eight) queues from BE (0) to NC (7). The packets destined for the CPU are classified internally and are put into the correct queue.

These packets are rate-limited to prevent DoS attacks. The software programs the classification entries to identify these packets and assigns appropriate bandwidth and priority to them. It is not configurable by the user.

2.10 Schedulers

This section provides information about QoS schedulers.

2.10.1 Scheduler modes on 7210 SAS-T

The scheduling modes interact with the minimum and maximum bandwidth CoS queue and maximum bandwidth egress port shaping specifications. Each egress port may be configured to have a specific scheduling mode. The scheduler first services the queues to meet their CIR and then services the queues to meet the PIR. There are five possibilities as follows:

- **Strict priority scheduling across CoS queues**

The strict priority scheduler provides strict priority access to the egress port across the CoS queue from highest CoS queue index (7) to the lowest (0). The purpose of the strict priority scheduler is to provide lower latency service to the higher CoS classes of traffic. In this mode, the scheduler services the queues in order of their priority in both the CIR and PIR loop.

As described in the following table, CoS queues 7 and 6 each have a minimum bandwidth specification of 10 Mbps, whereas the remaining QoS queues have a minimum bandwidth specification of 50 Mbps. All CoS queues have a maximum bandwidth specification of 1 Gbps.

The goal of these settings is to guarantee the minimum bandwidth settings for each of the queues while also allowing each CoS queue to fully use the egress port capability by having the maximum bandwidth setting at 1 Gbps. The strict priority scheduler mode provides low latency service for CoS queues 6 and 7 while their minimum bandwidth guarantees are being satisfied.

Table 45: Minimum and maximum bandwidth meters example

QoS queue name	Minimum bandwidth	Maximum bandwidth
7	10 Mbps	1 Gbps
6	10 Mbps	1 Gbps
5	50 Mbps	1 Gbps
4	50 Mbps	1 Gbps
3	50 Mbps	1 Gbps
2	50 Mbps	1 Gbps
1	50 Mbps	1 Gbps
0	50 Mbps	1 Gbps

- **Round robin scheduling across CoS queues**

The round robin scheduler mode provides round robin arbitration across the CoS queues. The scheduler visits each backlogged CoS queue, servicing a single packet at each queue before moving on to the next one. The purpose of the round robin scheduler is to provide fair access to the egress port bandwidth (at a packet level). This works best when packet sizes are approximately comparable. In this mode, the scheduler services the queues in round-robin for both the CIR and the PIR loop.

- **Weighted round robin (WRR)**

In WRR mode, the scheduler provides access to each CoS queue in round robin order. When the scheduler is providing access to a particular queue, it services a configurable number of back-to-back packets before moving on to the subsequent CoS queue. A value of strict is used to designate that a particular queue be considered to be a part of a hybrid Strict + WRR configuration.

The values 1 to 15 are used to indicate the number of back-to-back packets to be serviced when the scheduler is servicing a particular CoS queue. If the weight specified is N, but if the number of packets in the queue is lesser than N, the scheduler continues working and moves on to the next backlogged queue. In this mode, with no strict queues configured, the scheduler services the queues in round robin in the CIR loop. The configured weights are not considered in the CIR loop. The weights are used only in the PIR loop.

- **Weighted deficit round robin (WDRR) scheduling**

An inherent limitation of the WRR mode is that bandwidth is allocated in terms of packets. WRR works well if the average packet size for each CoS queue flow is known. WDRR aims at addressing this issue. WDRR provides a bandwidth allocation scheduler mode that takes into account the variably-sized packet issue by maintaining sufficient state information when arbitrating across the CoS queues. In this mode, with no strict queues configured, the scheduler services the queues in round-robin in the CIR loop.

The configured weights are not considered in the CIR loop. The weights are used only in the PIR loop. A weight value of 1 to 15 can be configured for each queue. Based on the weights provided respective amount of bytes is de-queued from the queue. A value of 0 is used to designate that a particular queue be considered to be a part of a hybrid Strict + WDRR configuration. If a weight value of 1 is given for queue 1 and 5 is given for queue 2, we see traffic out of the port in the ratio of 1:5 between the queues

(1 and 2), provided no traffic is flowing in the other queues. A weight value of 1 will actually pump out 2Kbytes from that queue, a value of 5 will pump out 10 Kbytes. Twice of the weight value given will be pumped out.

- **Strict + WRR/WDRR**

If the WRR/WDRR weight associated with a particular CoS queue is set to strict, the queue is considered to be in a strict priority mode. This set of strict priority queues is serviced first in the order of their CoS numbering (higher numbered CoS queue receives service before smaller numbered queues). In this mode, the scheduler services the strict queues first and then the queues configured with weights in both the CIR and PIR loop. The scheduler ensures that it meets the CIR of all the queues (both strict queues and queues with weight), if bandwidth is available before scheduling the queues in the PIR loop. If multiple queues are configured as strict, the higher-priority strict queues are serviced first before the lower priority strict queues in both the CIR and the PIR loop. The weights configured for the queues are only considered during the PIR loop.

2.10.2 Port scheduler policies for 7210 SAS-T

Port scheduler policies control the traffic behavior on a per-port basis. Associated with each egress port is a set of 8 (eight) class of service (CoS) queues and a default port-scheduler-policy named "default". This default policy makes the port to behave in strict mode. The default policy cannot be modified. The user can attach another policy to the port to change its scheduling behavior.

The scheduler that provides the arbitration across the 8 (eight) CoS queues is a scheduler that is configured in a variety of modes. A major aspect of the arbitration mechanism is the ability to provide minimum and maximum bandwidth guarantees. This is accomplished by tightly integrating a network queue and access egress policies into the scheduler. After the packets are mapped into a CoS queue, they are forwarded/conditioned using one of these schedulers (such as Strict Priority (SP), Round-Robin (RR), Weighted Round-Robin (WRR), Weighted Deficit Round-Robin (WDRR), (WRR+SP, WDRR+SP). The traffic shaping aspect is tightly integrated with the scheduler.

2.10.3 Schedulers on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE

The 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE support port-based scheduling with the following:

- per-port egress scheduler for access port
- per-port egress scheduler for network and hybrid ports

See [Schedulers on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#) for more information about scheduler behavior and to understand the queue parameters considered by the scheduler.

2.10.4 Schedulers on 7210 SAS-Mxp

7210 SAS-Mxp supports scheduling as follows:

- When SAP-based scheduling mode is enabled, the following support is available:
 - per-SAP egress scheduler for access port and hybrid port.
 - per-port egress scheduler for network and hybrid port.
- When port-based scheduling mode is enabled, the following support is available:

- per-port egress scheduler for access port.
- per-port egress scheduler for network and hybrid port.

See the [Schedulers on 7210 SAS-Mxp](#) chapter for more information about scheduling behavior and to understand the queue parameters considered by the scheduler.

2.10.5 Schedulers on 7210 SAS-R6 and 7210 SAS-R12

The 7210 SAS-R6 and 7210 SAS-R12 support scheduling as follows.

- When SAP-based scheduling mode is enabled, the following support is available:
 - per-SAP egress scheduler for access port and hybrid port
 - per-port egress scheduler for network and hybrid port
- When port-based scheduling mode is enabled, the following support is available:
 - per-port egress scheduler for access port
 - per-port egress scheduler for network and hybrid port

See [Schedulers on 7210 SAS-R6 and 7210 SAS-R12](#) for more information about scheduling behavior and queue parameters considered by the scheduler.

2.11 Configuration notes

The following information describes QoS implementation restrictions:

- Creating additional QoS policies is optional.
- Default policies are created for service ingress, access service egress, network, network-queue, slope, and port scheduler.
- Associating a service or access uplink or IP interface or network ports with a QoS policy other than the default policy is optional.

3 Discard eligibility indicator-based (DEI-based) classification and marking

This section provides information about the Discard Eligibility Indicator (DEI) feature that describes the requirements for DEI-based classification and marking for 7210 SAS platforms.

3.1 DEI-based classification

DEI-based classification is supported on access ports, access-uplink ports, network ports and hybrid ports as applicable on 7210 SAS platforms. DEI bits in the received packet are used to determine the ingress profile for the packet. If DEI = 0 in the received packet then the packet is considered to be GREEN or in-profile, and if DEI = 1 then the packet is considered to be YELLOW or out-of-profile. The profile assigned at the ingress can be used to enable color-aware metering with SAP ingress policing, network port ingress policing and access-uplink port ingress policing.

The profile of the packet can be reassigned by ingress meters/policers, when policing is used on SAP ingress, the final profile of the packet is determined by the meter/policers, based on the configured CIR/PIR rates. If a packet is below the CIR rate, it is assigned green/in-profile and if it exceeds the CIR rate and is below the PIR rate, it is assigned yellow/out-of-profile.

On the 7210 SAS, the behavior is the same when using ingress policing but is different when using ingress queuing. When using SAP ingress queuing, the profile assigned to the packet by user configuration cannot be reassigned by the ingress meters/policers or by ingress queue rate shapers. Therefore, the user-assigned profile is the final profile assigned to the packet.

The final profile assigned at ingress is used by egress to determine the WRED slope to use. The WRED slope determines whether the packet is eligible to be assigned a buffer and can be queued up on egress queue for transmission.



Note:

Ingress policing is supported on all 7210 SAS platforms as described in this document.

The following support is available for DEI classification:

- Under the port configuration, a command is provided to enable DEI-based classification, allowing user an option to enable/disable use of DEI for ingress classification on a per port basis. The initial profile (also known as color) is based on the DEI/CFI bit. If DEI = 0 in the received packet, then the packet can be considered GREEN or in-profile and if DEI = 1, then packet can be considered YELLOW or out-of-profile by the subsequent processing flow in hardware.
- All the SAPs configured on the port (access or hybrid) can use DEI classification for color-aware metering if required. The user has an option to use color-blind metering for some SAPs and color-aware metering for some other SAPs configured on the same port when DEI classification is enabled on the port. When using color-blind mode, the ingress profile assigned to the packet based on the DEI bit is ignored.
- The user is provided with an option in the SAP ingress QoS policy, to configure a policer as color-aware or color-blind. In color-aware mode, the DEI bit in the packet determines the ingress profile of the

packet. If the user configures the meter/policer mode as color-aware, the DEI bit of the incoming packet is used by the policer as the ingress profile.

- When using policing, the final profile of the packet is assigned by the ingress meter (based on configured CIR/PIR rate) in both color-aware and color-blind mode.
- On hybrid ports, the software allows only ONE of the following options to be configured:
 - If DEI-based classification is enabled, network port ingress policy **MUST** use Dot1p classification criteria with DEI profile for all configured Dot1p values.
 - OR
 - If DEI-based classification is disabled, network port ingress policy can use Dot1p or DSCP classification criteria.
- For network port policies, DEI-based classification is supported only when Dot1p classification criteria is in use. In other words, DEI-based classification cannot be used when DSCP based classification is used.
- For network IP interface policies DEI-based classification is not supported. In other words, when using EXP-based classification, DEI bit cannot be used to assign the profile for the packet.

3.2 DEI-based marking

DEI-based marking is supported on access ports, access-uplink ports, network ports, and hybrid ports. DEI bits can be used to mark the packet to carry the profile, assigned by an operator's trusted node at the ingress to the carrier's network, to the subsequent nodes in the network. It allows high-priority in-profile packet to be allocated appropriate resources by all the network nodes on the path to the final destination. Similarly, it allows out-of-profile packets to be treated with less preference compared to in-profile packets by all the network nodes on the path to the final destination. The egress marking behavior must be symmetric to the ingress classification behavior.

The following support is available for DEI-based marking:

- Option to mark DEI bits for access SAP egress on access ports, network ports, and hybrid ports on 7210 SAS devices configured in network mode.
- Option to mark DEI bits for port egress on access ports and access-uplink ports on 7210 SAS devices configured in access-uplink mode.
- Option to mark DEI bits for IP and MPLS packets on network ports (DEI marking is supported for MPLS packets only on those platforms that support dot1p marking for MPLS packets. DEI marking is not supported otherwise).
- By default, in-profile packets are marked with DEI bit = 0 and out-of-profile packets are marked with DEI bit = 1. The user has an option to mark all the packets belonging to a FC to the same DEI value irrespective of its profile using the **force de-mark** command.



Note:

For information about the commands for configuring DEI, see [Network QoS policies](#), [Network queue QoS policy command reference](#), [Access egress QoS policies on 7210 SAS-T, 7210 SAS-Sx/S 1/10GE](#), and [7210 SAS-Sx 10/100GE](#), [Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#), [Access egress QoS policy command reference](#), and the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide*.

3.3 Configuration guidelines

The following are configuration guidelines for DEI-based classification and marking:

- While disabling DEI-based classification on a port, all the meters used by the SAPs configured on this port must be in color-blind mode. The converse is also true; that is, while attaching a SAP ingress QoS policy with meter as color-aware to a SAP, the DEI-based classification must have been enabled on the port on which SAP exists.
- While configuring DEI-based classification in a network port ingress policy, only Dot1p classification can be used.
- DEI classification must be disabled on that port before changing the mode from one mode (access/network/hybrid) to another mode.
- All the ports under a LAG should have the same configuration for DEI classification (enable/disable). If the LAG configuration changes, the port configuration also will be updated accordingly. Port configuration under the LAG cannot be changed.
- While enabling DEI-based policing on a port, if it is a hybrid port, then the network port ingress QoS policy must use only Dot1p mappings for classification and all the configured profiles must be **use-dei** (not **in** or **out**). This is true even if only SAPs are configured on the port.



Note:

- Only after attaching a network port ingress QoS policy as mentioned above, can users enable DEI classification on a port.
 - For more information about DEI classification on a LAG and port, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide*.
- While configuring the profile on a network port policy for Dot1p classification, if the policy is attached to a hybrid port on which DEI classification is enabled, the profile cannot be changed to **in** or **out**. In other words, only DEI bits can be used for profile configuration.
 - While attaching a network QoS policy to a hybrid port, all the Dot1p mappings including the default mapping should use only the DEI as profile, if DEI classification is enabled on this port.



Note:

- DEI-based classification cannot be configured for IP interfaces.
- Only MPLS EXP-based classification is available for IP interfaces.

4 Port level egress rate-limiting

This section provides information to configure port level **egress-rate** command using the command line interface.

4.1 Overview

Egress port rate limiting allows the device to limit the traffic that egresses through a port to a value less than the available link bandwidth. This feature is supported on the 7210 SAS-series platforms.

4.1.1 Applications

This feature is useful when connecting the 7210 SAS to an Ethernet-over-SDH (EoSDH) (or microwave) network, where the network allocates predetermined bandwidth to the nodes connecting into it, based on the transport bandwidth requirement. When connecting to such a network it is important that the traffic sent into the SDH node does not exceed the configured values because the SDH network does not have QoS capabilities and buffers required to prioritize the ingress traffic.

Egress rate attributes include:

- Allows for per port configuration of the maximum egress port rate, using the **egress-rate** CLI command.
- Ethernet ports configured as access, access uplink and network support this feature.
- The port scheduler distributes the available maximum egress bandwidth based on the CIR/PIR configuration parameters provisioned for the queues.
- Provides support for a burst parameter to control the amount of burst the egress port can generate.
- When ports are members of a LAG, all the ports use the same value for the **egress-rate** and the **max-burst** parameters.
- If frame overhead accounting is enabled, then queue scheduler accounts for the Ethernet frame overhead.

4.1.2 Effect of port level rate-limiting on network queue functionality

- When an **egress-rate** sub-rate value is given, the network queue (on network ports or access uplink ports) rates that are specified using percentages will use the **egress-rate** value instead of the port bandwidth (if egress rate is lesser than port bandwidth) to configure the appropriate queue rates. Configuration of egress port rate to different values will result in a corresponding dynamic adjustment of rates for the queues configured on network ports, or access uplink ports.
- When the **egress-rate sub-rate** value is set, CBS/MBS of the associated network queues will not change.

4.2 Basic configurations

To apply port level rate-limiting, perform the following:

- The **egress-rate** command is present in the ***A:Dut-1>config>port>ethernet** context.
- The **egress-rate** configures the maximum rate (in kbps) for the port. The value should be between 1 and 1000000 kbps and between 1 and 10000000 kbps for a 10G port.
- The **max-burst** command configures a maximum-burst (in kilo-bits) associated with the egress-rate. This is an optional parameter and if not defined then, by default, it is set to 64 kb for a 1G port and 64 kb for a 10G port. Users cannot configure **max-burst** without configuring **egress-rate**. The **max-burst** value should be between 32 and 16384 or the default.
- By default there is no **egress-rate** command set on port. The default **egress-rate** for a port is the maximum (equal to line-rate).
- On 10G port, if Egress port Rate Limiter (ERL) configured is more than 8Gig, Nokia recommends configuring the burst value higher than 80kbits to avoid packet drops.

Output example

The following is a sample configuration output that shows the **egress-rate** configuration for a port.

```
*A:Dut-1>config>port# info
-----
    ethernet
        egress-rate 120000 max-burst 234
    exit
    no shutdown
-----
*A:Dut-1>config>port#
```

4.2.1 Modifying port level egress-rate command

To modify egress rate parameters you can simply apply an **egress-rate** command with new **egress-rate** and **max-burst** value.

Output example

The following is a sample configuration that shows a modified **egress-rate** configuration for a port.

```
*A:Dut-1>config>port# ethernet egress-rate 10000 max-burst default
*A:Dut-1>config>port# info
-----
    ethernet
        egress-rate 10000
    exit
    no shutdown
-----
*A:Dut-1>config>port#
```

4.2.2 Removing port level egress-rate command

To remove **egress-rate** command from a port, use the **no** option with the **egress-rate** command. Do not include the rate for the **egress-rate** and **max-burst** options. Use the following syntax to remove the **egress-rate** command from a port.

```
config>port>ethernet# no egress-rate
```

Output example

The following is a sample configuration output that shows the removal of **egress-rate** configuration from a port.

```
*A:Dut-1>config>port# no ethernet egress-rate
*A:Dut-1>config>port# info
-----
      ethernet
      exit
      no shutdown
-----
*A:Dut-1>config>port#
```

4.2.2.1 Default egress-rate values

By default no **egress-rate** is configured for a port.

4.3 Port level egress-rate command reference

4.3.1 Command hierarchies

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

4.3.1.1 Configuration commands

```
config
- port
  - ethernet
    - egress-rate sub-rate [max-burst size-in-kbits]
    - no egress-rate
```

4.3.1.2 Show commands

```
show
- port [port-id]
```

4.3.2 Command descriptions

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

4.3.2.1 Configuration commands

egress-rate

Syntax

egress-rate *sub-rate* [**max-burst** *size-in-kbits*]
no egress-rate

Context

config>port>ethernet

Platforms

Supported on all 7210 SAS platforms as described in this document, including platforms configured in the access-uplink operating mode

Description

This command configures the maximum rate and corresponding burst value for a port.
The **no** form of this command removes the **egress-rate** from the port.

Parameters

- sub-rate**

Specifies the maximum rate, in kbps.

Values	1 to 1000000
	1 to 10000000 (10G port)
- max-burst size-in-kbits**

Specifies the maximum burst size, in kbs.

Values	32 to 16384, default
Default	64

4.3.2.2 Show commands

```
port
```

Syntax

```
port [port-id]
```

Context

```
show
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including platforms configured in the access-uplink operating mode

Description

This command displays the egress rate and max burst value set for the port, along with other details about the port.

Parameters

port-id
Displays information about the specific port ID.

Output

The following output is an example of port information, and [Table 46: Output fields: specific port](#) describes the output fields.

Sample output

```
*A:Dut-1>config>port>ethernet# show port 1/1/23
=====
Ethernet Interface
=====
Description      : 10/100/Gig Ethernet SFP
Interface        : 1/1/23
Link-level       : Ethernet
Admin State      : up
Oper State       : up
Physical Link    : Yes
IfIndex          : 36405248
Last State Change : 03/12/2001 03:31:09
Last Cleared Time : N/A
Oper Speed       : 100 mbps
Config Speed     : 1 Gbps
Oper Duplex      : full
Config Duplex    : full
MTU              : 9212
Hold time up    : 0 seconds
Hold time down  : 0 seconds

Configured Mode  : network
Dot1Q Ethertype : 0x8100
Net. Egr. Queue Pol: default
Egr. Sched. Pol : default
Auto-negotiate  : true
Accounting Policy : None
Egress Rate     : 100000
Encap Type      : null
QinQ Ethertype  : 0x8100
Access Egr. Qos *: n/a
Network Qos Pol : 1
MDI/MDX         : MDX
Collect-stats   : Disabled
Max Burst       : 8000
```



```

Down-when-looped : Disabled           Keep-alive       : 10
Loop Detected    : False              Retry           : 120

Configured Address : 00:f7:d6:5e:98:18
Hardware Address   : 00:f7:d6:5e:98:18
Cfg Alarm         :
Alarm Status      :

Transceiver Data

Transceiver Type   : SFP
Model Number       : 3HE00062AAAA01 ALA IPUIAEHDAA6
TX Laser Wavelength: 0 nm              Diag Capable    : no
Connector Code     : Unknown            Vendor OUI      : 00:90:65
Manufacture date    : 2008/09/11         Media           : Ethernet
Serial Number       : PEB2WGH
Part Number        : FCMJ-8521-3-A5
Optical Compliance : GIGE-T
Link Length support: 100m for copper

=====
Traffic Statistics
=====
                                Input           Output
-----
Octets                  15028477             3236
Packets                 16729                19
Errors                   0                  0
=====
* indicates that the corresponding row element may have been truncated.

=====
Port Statistics
=====
                                Input           Output
-----
Unicast Packets         11611             17
Multicast Packets        359                0
Broadcast Packets       4759                2
Discards                0                  0
Unknown Proto Discards  0
=====

=====
Ethernet-like Medium Statistics
=====

Alignment Errors :      0 Sngl Collisions :      0
FCS Errors       :      0 Mult Collisions :      0
SQE Test Errors  :      0 Late Collisions :      0
CSE              :      0 Excess Collisns :      0
Too long Frames  :      0 Int MAC Tx Errs :      0
Symbol Errors    :      0 Int MAC Rx Errs :      0
=====
*A:MTU-T2>config>port>ethernet#

```

Table 46: Output fields: specific port

Label	Description
Description	A text description of the port

Label	Description
Interface	The physical port ID in the form <i>slot/mda/port</i> for non-QSFP28 ports, or <i>slot/mda/cport/channel</i> for QSFP28 ports
Oper Speed	The speed of the interface
Link-level	Ethernet — The port is configured as Ethernet
Config Speed	The configured speed of the interface
Admin State	Up — The port is administratively up Down — The port is administratively down
Oper Duplex	Full — The link is set to full duplex mode Half — The link is set to half duplex mode
Oper State	Up — The port is operationally up Down — The port is operationally down Additionally, the <i>lag-id</i> of the LAG it belongs to in addition to the status of the LAG member (active or standby) is specified.
Config Duplex	Full — The link is set to full duplex mode Half — The link is set to half duplex mode
Physical Link	Yes — A physical link is present No — A physical link is not present
MTU	The size of the largest packet which can be sent/received on the Ethernet physical interface, specified in octets
IfIndex	Displays the interface's index number which reflects its initialization sequence
Hold time up	The link up dampening time in seconds The port link dampening timer value which reduces the number of link transitions reported to upper layer protocols.
Last State chg	Displays the system time moment that the peer is up
Hold Time Down	The link down dampening time in seconds The down timer controls the dampening timer for link down transitions.
Last Cleared Time	The time since the last clear
Configured Mode	network — The port is configured for transport network use access — The port is configured for service access
Encap Type	Null — Ingress frames will not use any tags or labels to delineate a service

Label	Description
	dot1q — Ingress frames carry 802.1Q tags where each tag signifies a different service QinQ — Encapsulation type specified for QinQ Access SAPs
Dot1q Ethertype	Indicates the Ethertype expected when the port's encapsulation type is dot1q
QinQ Ethertype	Indicates the Ethertype expected when the port's encapsulation type is QinQ
Net. Egr. Queue Pol	Specifies the network egress queue policy or that the default policy is used
Access Egr. Qos	Specifies the access egress policy or that the default policy 1 is in use
Egr. Sched. Pol	Specifies the port scheduler policy or that the default policy default is in use
Network Qos Pol	The network QoS policy ID applied to the port
Auto-negotiate	True — The link attempts to automatically negotiate the link speed and duplex parameters False — The duplex and speed values are used for the link
MDI/MDX	Indicates the Ethernet interface type
Accounting Policy	Indicates the accounting policy, if configured
Collect-stats	Enabled — The collection of accounting and statistical data for the network Ethernet port is enabled When applying accounting policies the data by default is collected in the appropriate records and written to the designated billing file. Disabled — Collection is disabled Statistics are still accumulated by the IOM cards, however, the CPU does not obtain the results and write them to the billing file.
Egress Rate	The maximum amount of egress bandwidth (in kilobits per second) that this Ethernet interface can generate
Max Burst	The maximum number of RSVP messages to be sent
Down-when-looped	Shows whether the feature is enabled or disabled
Keep-alive	The interval at which keep-alive messages are exchanged
Loop Detected	Indicates whether a loop is detected on the port

Label	Description
Retry	Indicates the minimum wait time before the port is re-enabled after it is brought down because of a loop detection
Configured Address	The base chassis Ethernet MAC address
Hardware Address	The interface's hardware or system assigned MAC address at its protocol sublayer
Transceiver Data	
Transceiver Type	Type of the transceiver
Model Number	The model number of the transceiver
Transceiver Code	The code for the transmission media
TX Laser Wavelength	The light wavelength transmitted by the transceiver laser
Connector Code	The vendor organizationally unique identifier field (OUI) contains the IEEE company identifier for the vendor
Diag Capable	Indicates if the transceiver is capable of doing diagnostics
Vendor OUI	The vendor-specific identifier field (OUI) contains the IEEE company identifier for the vendor
Manufacture date	The manufacturing date of the hardware component in the mmddyyyy ASCII format
Media	The media supported for the SFP
Serial Number	The vendor serial number of the hardware component
Part Number	The vendor part number contains ASCII characters, defining the vendor part number or product name
Traffic Statistics	
Input/Output	When the collection of accounting and statistical data is enabled, octet, packet, and error statistics are displayed
Octets	Total number of octets received
Packets	The number of packets received, broken down by size Port Statistics
Port Statistics	
Errors Input/Output	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol

Label	Description
	<p>For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol</p> <p>For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors</p> <p>For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors</p>
Unicast Packets Input/Output	<p>The number of packets, delivered by this sublayer to a higher sublayer, which were not addressed to a multicast or broadcast address at this sublayer</p> <p>The total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a multicast or broadcast address at this sublayer, including those that were discarded or not sent</p>
Multicast Packets Input/Output	<p>The number of packets, delivered by this sublayer to a higher sublayer, which were addressed to a multicast address at this sublayer</p> <p>For a MAC layer protocol, this includes both group and functional addresses.</p> <p>The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sublayer, including those that were discarded or not sent</p> <p>For a MAC layer protocol, this includes both Group and Functional addresses.</p>
Broadcast Packets Input/Output	<p>The number of packets, delivered by this sublayer to a higher sublayer, which were addressed to a broadcast address at this sublayer</p> <p>The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sublayer, including those that were discarded or not sent</p> <p>For a MAC layer protocol, this includes both Group and Functional addresses.</p>
Discards Input/Output	<p>The number of inbound packets chosen to be discarded to possibly free up buffer space</p>
Unknown Proto Discards Input/Output	<p>For packet-oriented interfaces, the number of packets received through the interface which were discarded because of an unknown or unsupported protocol</p> <p>For character-oriented or fixed-length interfaces that support protocol multiplexing the number of transmission units received</p>

Label	Description
	via the interface which were discarded because of an unknown or unsupported protocol For any interface that does not support protocol multiplexing, this counter is always 0.
Ethernet-like Medium Statistics	
Alignment Errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but that had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets
FCS Errors	The number of frames received that are an integral number of octets in length but do not pass the FCS check
SQE Errors	The number of times that the SQE TEST ERROR is received
CSE	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame
Too long Frames	The number of frames received that exceed the maximum permitted frame size
Symbol Errors	For an interface operating at 100 Mb/s, the number of times there was an invalid data symbol when a valid carrier was present
Sngl Collisions	The number of frames that are involved in a single collision, and are subsequently transmitted successfully
Mult Collisions	The number of frames that are involved in more than one collision and are subsequently transmitted successfully
Late Collisions	The number of times that a collision is detected later than one slot Time into the transmission of a packet
Excess Collisions	The number of frames for which a transmission fails because of excessive collisions
Int MAC Tx Errs	The number of frames for which a transmission fails because of an internal MAC sublayer transmit error
Int MAC Rx Errs	The number of frames for which a reception fails because of an internal MAC sublayer receive error

5 SAP egress aggregate meter

This chapter provides information to configure service-level egress rate limits using the command line interface.

5.1 Overview

The SAP egress aggregated meter feature allows a user to limit the amount of traffic sent out of a SAP by enforcing service-level egress rate limits.

This feature is supported only on the 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (network and access-uplink mode), 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE platforms. This feature is not supported on 7210 SAS-Mxp.

This feature allows a meter to be associated with the SAP egress to limit the aggregate amount of traffic across all FCs that are sent out of the SAP. This feature is supported for SAPs in P2MP services (Epipe, VPLS, RVPLS, IES, and VPRN).

5.1.1 Configuration notes

- On the 7210 SAS-R6 and 7210 SAS-R12, if using per-SAP egress queuing instead of port-based egress queuing, the per-SAP egress aggregate shaper rate can be used instead of this feature.
- Before configuring a SAP-egress aggregate meter on a SAP, the user must reallocate the resources from the egress-internal-tcam pool toward the SAP egress meters using the **configure>system>resource-profile>egress-internal-tcam>egress-sap-aggregate-meter** CLI command. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* or details about the **egress-sap-aggregate-meter** command. The **egress-internal-tcam** pool resources are shared with other features, such as egress ACLs.
- The SAP egress aggregate policer is not FC and profile aware. The packets scheduled out of the per-port queues are on a first-come-first-served basis and limits the amount of traffic to the configured rate across all the FCs for a SAP.
- See the *7210 SAS-R6, R12 Services Guide* or *7210 SAS-Mxp, S, Sx, T Services Guide* for details about the optional **enable-stats** command under the **configure service vpls sap egress aggregate-meter-rate** CLI command. This command associates a counter with the meter to count the forwarded and dropped packets and octets. If the command is enabled, and accounting is enabled for the SAP, service egress accounting records contain the forwarded and dropped counts.

5.2 Basic configurations

About this task

To enable the per-SAP egress aggregate meter command, perform the following steps.



Note:

Step 1 applies only to the 7210 SAS-R6 and 7210 SAS-R12.

Procedure

Step 1. Ensure that **port-scheduling-mode** is enabled on the node by doing the following:

- Use the **show system global-res-profile active** command to display the current scheduling mode for the node. The **port-scheduler-mode** should be displayed as "enable".
- If the port-scheduler-mode is "disable", use the **configure system global-res-profile qos port-scheduler-mode** command to enable it.

The following output displays the port-scheduler-mode status.

```
=====
Active Global System Resource Profile Information
=====
-----
port-scheduler-mode      : enable
-----
```

Step 2. Allocate resources to the feature using the **configure system resource-profile egress-internal-tcam egress-sap-aggregate-meter** command. If necessary, free up resources for use by this feature by removing resources from other features.

Step 3. Configure the SAP-egress aggregate meter rate using the **configure service vpls sap egress aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] [enable-stats]** command



Note:

Use the **enable-stats** parameter to enable the counter that is used to count total forwarded packets out of the SAP.

5.3 SAP egress aggregate meter command reference

See the *7210 SAS-Mxp, S, Sx, T Services Guide* or *7210 SAS-R6, R12 Services Guide* for information about the commands described in this section.

6 Frame-based accounting

This chapter provides information to configure frame-based accounting using the command line interface.

6.1 Overview

This feature when enabled allows QoS policies to account for the Ethernet frame overhead (for example, it accounts for the IFG (inter-frame gap) and the preamble). Typically, the IFG and preamble constitutes about $12 + 8 = 20$ bytes. The overhead for Ethernet ports uses this value.

6.1.1 Frame-based accounting

A configurable CLI command enables accounting of the frame overhead at ingress or egress. This is a system wide parameter and affects the behavior of the ingress meter or egress rate. When disabled, the queue rates and egress-rate do not account for the Ethernet frame overhead. By default frame-based accounting is disabled for both ingress and egress. Frame overhead is always accounted for on egress (queue rates and egress rate) and the user has no option to disable it.



Note:

- The egress port rate limiter (ERL) calculation is also frame based, and the user has no option to disable it; ERL can be applied on the port to shape the egress rate.
- On the 7210 SAS-Mxp, frame-based accounting for service meters is not supported.
- Frame-based accounting for SAP egress aggregate meters is not supported on the 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-R6 and 7210 SAS-R12 in port scheduler mode.

6.1.2 Effects of enabling ingress frame-based accounting on ingress meter functionality

To enable system-wide consistency in configuring QoS queue and meter rate parameters, the meters used on the system ingress may need to account for Ethernet frame overhead. Network ingress and service ingress meters account for Ethernet frame overhead. A configurable CLI command can enable or disable the frame overhead accounting. This is a system-wide parameter affecting the behavior of all the meters in the system.

6.1.3 Effects of enabling egress frame-based accounting on network queue functionality

If frame overhead consideration is enabled, then queue scheduler accounts for the Ethernet frame overhead. The maximum egress bandwidth accounts for the Ethernet frame overhead (it accounts for the IFG (inter-frame gap) and the preamble). Typically, the IFG and preamble constitutes about $12 + 8 = 20$ bytes. The overhead for Ethernet ports uses this value.

A configurable CLI command enables accounting of the frame overhead. This is a system wide parameter and affects the behavior of all egress queues (when **frame-based-accounting** is enabled on egress port (network ports or access-uplink ports, as applicable), the associated queues also account for frame overhead implicitly). When disabled, the **egress-rate** command does not account for the Ethernet frame overhead.

6.1.4 Accounting and statistics

Accounting records and statistics do not account for frame overhead.

6.2 Basic configurations

To enable **frame-based accounting**, you must perform the following:

- The **frame-based-accounting** command is in the ***A:Dut-1>config>qos>frame-based-accounting** context.
- The **ingress-enable** command enables **frame-based-accounting** for ingress metering.
- The **egress-enable** command enables **frame-based-accounting** for egress queue rates, scheduler and port level egress-rate.

Output example

The following is a sample frame-based accounting configuration output.

```
*A:Dut-1>config>qos>frame-based-accounting# info detail
-----
                no ingress-enable
                no egress-enable
-----
*A:Dut-1>config>qos>frame-based-accounting#
```

6.2.1 Enabling and disabling frame-based accounting

To enable **frame-based-accounting** for ingress, you can simply use the **ingress-enable** command and to enable **frame-based-accounting** on egress use the **egress-enable** command. To disable **frame-based-accounting** for ingress, execute the **no ingress-enable** command and to disable **frame-based-accounting** on egress, execute the **no egress-enable** command.

```
config>qos>frame-based-accounting
```

Output example

The following is a sample configuration output that shows the enabling of frame-based accounting.

```
*A:Dut-1>config>qos>frame-based-accounting# ingress-enable
*A:Dut-1>config>qos>frame-based-accounting# egress-enable
*A:Dut-1>config>qos>frame-based-accounting# info
-----
                ingress-enable
                egress-enable
-----
```

```
*A:Dut-1>config>qos>frame-based-accounting#
```

Output example

The following is a sample configuration output that shows the disabling of frame-based accounting.

```
*A:Dut-1>config>qos>frame-based-accounting# no ingress-enable
*A:Dut-1>config>qos>frame-based-accounting# no egress-enable
*A:Dut-1>config>qos>frame-based-accounting# info detail
-----
                no ingress-enable
                no egress-enable
-----
*A:Dut-1>config>qos>frame-based-accounting#
```

6.2.1.1 Default frame-based accounting values

By default, the **frame-based-accounting** command is disabled for ingress.

By default, the **frame-based-accounting** command is disabled for egress. It is user-configurable and can be enabled or disabled.

6.3 Frame-based accounting command reference

6.3.1 Command hierarchies

- [Configuration commands](#)

6.3.1.1 Configuration commands

```
config
- qos
  - frame-based-accounting
    - [no] egress-enable
    - [no] ingress-enable
```

6.3.2 Configuration commands

frame-based-accounting

Syntax

frame-based-accounting

Context

config>ops

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in the access-uplink mode

Description

Commands in this context configure frame-based accounting.

egress-enable

Syntax

[no] egress-enable

Context

config>qos>frame-based-accounting

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in the access-uplink mode

Description

This command enables the frame-based-accounting for access-egress, network-queue, port scheduler, SAP or Network Aggregate Rate and port-level egress-rate.

The **no** form of this command disables frame-based-accounting for all egress QoS.

Default

no egress-enable

ingress-enable

Syntax

[no] ingress-enable

Context

config>qos>frame-based-accounting

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in the access-uplink mode

Description

This command enables the frame-based accounting for SAP-ingress and network QoS.

The **no** form of this command disables frame-based accounting for SAP-ingress and network QoS.

Default

no ingress-enable

7 Network QoS policies

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

7.1 Overview

This section provides an overview of QoS policies in network mode and access up-link mode.

7.1.1 Overview of network QoS policies in network mode

Network QoS policies are available for use with network IP interfaces, network ports and hybrid ports when operating the 7210 SAS in network and standalone mode. The 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 support network and standalone mode of operation.

The following types of QoS mapping decisions are applicable on a network IP interface when operating in network and standalone mode:

- MPLS LSP EXP value mapping to FC (if defined)
- default QoS mapping
- MPLS LSP EXP mapping to profile

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

The following types of QoS mapping decisions are applicable on a network port when operating in network mode:

- Ethernet dot1P and IP DSCP value mapping (if defined) for use with IP packets
- default QoS mapping

The default QoS mapping always exists on network port and every received packet is mapped to the default if another explicitly defined matching entry does not exist.

7.1.1.1 Overview of network QoS policies in access-uplink mode

Network QoS policies are available for use with access-uplink port when operating the 7210 SAS in access-uplink mode. The 7210 SAS-T supports the access-uplink mode of operation.

The following types of QoS mapping decisions are applicable on an access-uplink port when operating in access-uplink mode:

- option to mark Ethernet dot1p or IP DSCP
- default QoS mapping

The default QoS mapping always exists on an ingress access uplink port and every received packet is mapped to this default if another explicitly defined matching entry does not exist.

The following sections provide more details about the network QoS policies for both network mode and access-uplink mode.

7.2 Network QoS policy in network mode

The network QoS policy consists of an ingress and egress component. For the 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 operating in network mode, there are two types of network QoS policies; network QoS policy of type **port** and network QoS policy of type **ip-interface**.

A **port** type network policy is applied to network and hybrid ports and is used for classification and remarking of IP traffic using DSCP or dot1p values. Either DSCP or dot1p can be used for ingress classification but not both. Both DSCP and dot1p can be configured at egress for remarking.

An **ip-interface** type network policy is applied to an IP Interface, and is used for classification and remarking of MPLS traffic using EXP values. Note that the FC-to-dot1p marking values configured on the port are also used to mark the dot1p in the VLAN tag, if any, used for MPLS traffic on some 7210 SAS devices.

The ingress component of the **port** policy type defines how the DSCP or dot1p bits are mapped to internal forwarding class (FC) and profile state. The ingress component of the **ip-interface** policy type defines how the EXP bits are mapped to the internal FC and profile state. The FC and profile state define the per-hop behavior (PHB) or the QoS treatment through the 7210 SAS. The ingress profile assignments using MPLS-EXP bits are defined using the **mpls-lsp-exp-profile-map** policy, which defines the mapping between the MPLS LSP EXP bits and the profile (in or out) associated with a packet.

The mapping on each **ip-interface** or **port** policy defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the IP interface or port. It also defines the bandwidth-limiting parameters for the traffic mapped to each FC. Traffic mapped to each FC can be limited to configurable bandwidth values using separate meters and queues for unicast and multipoint traffic. Multipoint is used for IP interfaces for MPLS traffic and for IP multicast traffic received on a network or hybrid port. On both network IP interface ingress and network port ingress, color-aware meters are provisioned by default to use the ingress profile assigned to the packet, if configured.

The total number of QoS resources—that is, ingress classification entries and policers—available for use with IP interfaces is limited. The software allocates these resources to an IP interface on a first-come, first-served basis. The number of resources used per IP interface limits the total number of IP interfaces configured on the system. The total number of IP interfaces allowed is also subject to a system limit.

The egress component of the network QoS **ip-interface** policy type defines the LSP EXP bits marking values associated with each FC. The egress component of the network QoS **port** policy type defines the DSCP or dot1p bits marking values associated with each FC.

By default, all ports configured in network mode use the default network policy "1" and all network port IP interfaces use the default network policy "2". Default network policies "1" and "2" cannot be modified or deleted.

New (non-default) network policy parameters can be modified. The **no** form of this command reverts to the default values.

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

See the “CLI Usage” chapter in the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for information about the tasks and commands required to access the CLI and to configure and maintain 7210 SAS devices.

7.2.1 Network QoS policy (ip-interface type) behaviour for MPLS LSPs

The following behavior is supported with use of network IP interface QoS policies for LDP, segment routing (SR), and RSVP (with fast reroute (FRR) or penultimate hop popping (PHP)) MPLS LSPs:

- The **mpls-lsp-exp-profile-map** policy allows you to specify the profile mapping and FC mapping independently. The FC to use is always taken from the network policy using the EXP-to-FC mapping configured in the network policy. The EXP-to-profile mapping defined in the **mpls-lsp-exp-profile-map** policy associated with the network QoS policy is used. Each IP interface can define a unique network policy for use, and each network policy uses a different MPLS LSP EXP-to-FC mapping. If adequate network classification resources are available, this mapping allows the use of more than 32 distinct network policies.
- MPLS transport tunnels that are set up using LDP and segment routing use a global **mpls-lsp-exp-profile-map** policy. Although the system assigns a default **mpls-lsp-exp-profile-map** policy, the user can change the default and specify the global policy to use.
- Use the **mpls-lsp-exp-profile-map** policy to assign profile values based on the MPLS EXP bits for MPLS packets received over different IP interfaces. This is particularly useful for RSVP LSPs with FRR 1:1 configurations.
- For LDP LSPs or when using the FRR facility, Nokia recommends that a single **mpls-lsp-exp-profile-map** policy should be used for all IP interfaces. If PHP is enabled, the egress LER receives only VC-labeled packets and the global **mpls-lsp-exp-profile-map** policy is used for profile mapping. Therefore, when PHP is enabled, a single **mpls-lsp-exp-profile-map** policy should be used for all IP interfaces to get consistent profile treatment.
- If traffic is received on a transport tunnel set up using RSVP LSP, LDP or SR with an identical EXP bit value, the system provides the same QoS treatment to the MPLS traffic. The FC and meter from the network QoS policy are used for all MPLS traffic received on an IP interface, regardless of whether it is LDP or RSVP LSP. The profile value assigned depends on the configuration. If the user needs similar profile mapping for MPLS traffic received across all transport tunnels, a single **mpls-lsp-exp-profile-map** must be used for all IP interfaces and for the global **mpls-lsp-exp-profile-map**.

7.2.2 Basic configurations

A basic network QoS policy must conform to the following:

- Each network QoS policy must have a unique policy ID.
- The network must specify the default action.
- The network must have a QoS policy scope of **template** or **exclusive**.

7.2.3 Create a network QoS policy (ip-interface type) for network mode

Configuring and applying QoS policies other than the default policy is optional. A default network policy of the type **ip-interface** is applied to each IP interface.

To create a network QoS policy of the **ip-interface** type when operating in network mode, define the following:

- Specify a network policy ID value. The system does not dynamically assign a value.
- Set the **network-policy-type** parameter to **ip-interface**.
- Include a description that provides a brief overview of policy features.
- Use egress marking and remarking to specify the egress LSP EXP marking map; otherwise, the default values are applied. The following are defined:

- **remarking**

The **remark** *policy-id* command specifies the policy, which defines the mapping of the FC-to-packet header priority profile and bits.

The **remarking** command uses the associated policy configured with the **remark** *policy-id* command to determine which priority bits to mark on egress.

When remarking is enabled, MPLS EXP bits for all MPLS LSR and LER traffic are marked on egress on the specified network IP interface. Remarking is based on the FC-to-LSP EXP bit mapping that is defined in the remark policy and associated under the egress node of the network QoS policy. The EXP marking values used depend on the configured remark policy. If the user does not attach an explicit policy, the default policy is used.

When remarking is disabled for MPLS LSR traffic, EXP values received on ingress are not modified on egress. For MPLS LER traffic where the node adds the MPLS encapsulation, MPLS EXP bits are set based on the mapping specified in the policy associated with the IP interface. If the user does not attach an explicit policy, the default policy is used.

- **FC criteria**

The FC name represents an egress queue. Specify FC criteria to define the egress characteristics of the queue and the marking criteria of packets flowing through it.

- **LSP EXP**

The EXP is used for all packets requiring marking that egress on this FC queue that are in or out of profile.

- Specify ingress criteria using the following:

- **default action**

Defines the default action to take for packets that have undefined bits set. The default action specifies the FC to which these packets are assigned.

- **LSP EXP**

Creates a mapping between the LSP EXP bits of the network ingress MPLS traffic and the FC. Ingress MPLS traffic that matches the specified EXP bits are assigned to the corresponding FC. The user has an option to specify the mapping of the LSP EXP bits to a profile (in or out). Ingress traffic that matches the specified EXP bits is assigned the corresponding profile.

Use the following syntax to create a network QoS policy.

```
config>router
  interface interface-name
    qos network-policy-id
```

7.2.3.1 Example for network QoS policy of ip-interface type

The following is a sample configuration output of a network **ip-interface** type QoS policy for the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T.

Output example

```
*7210 SAS>config>qos>network# info
-----
      ingress
        meter 1 create
        exit
        meter 3 create
        exit
        meter 5 multipoint create
        exit
        meter 9 multipoint create
        exit
        fc "ef" create
          meter 3
          multicast-meter 5
        exit
        lsp-exp 0 fc be
        lsp-exp 1 fc l2
        lsp-exp 2 fc af
        lsp-exp 3 fc af
        lsp-exp 4 fc h2
        lsp-exp 5 fc ef
      exit
    egress
      remarking
      remark 200
    exit
-----
```

7.2.4 Configuring network QoS policy (port type) for network mode

To create a network QoS policy of the **port** type when operating in network mode, define the following:

- Specify a network policy ID value. The system does not dynamically assign a value.
- Set the **network-policy-type** parameter to **port**.
- Include a description that provides a brief overview of policy features.
- Modify the egress DSCP and dot1p marking map; otherwise, the default values are applied:

- **remarking**

When enabled, this command remarks all IP packets that egress the specified network port. The remarking is based on the FC-to-DSCP bit mapping defined in the remark policy and associated under the egress node of the network QoS policy for all IP traffic. FC-to-dot1p marking for MPLS packets is not supported; it is supported only for IP packets.

- **FC criteria**

The FC name represents an egress queue. Specify FC criteria to define the egress characteristics of the queue and the marking criteria of packets flowing through it.

- **DSCP and dot1p**

Specify the DSCP and dot1p value to use for IP packets requiring marking that egress on this FC queue that are in or out of profile.

- Specify ingress criteria as either DSCP or dot1p (but not both) to FC mapping for all IP packets. Define the following:
 - **default action**
Defines the default action to take for packets that have undefined DSCP or dot1p bits set. The default action specifies the FC to which these packets are assigned.
 - **DSCP or dot1p**
Creates a mapping between the DSCP or dot1p bits of network ingress IP traffic and the FC. Ingress IP traffic that match the specified DSCP or dot1p bits are assigned to the corresponding FC.

The following example shows the command usage to associate a network QoS policy with the network port.

```
network port (in network mode)
config> port
    ethernet
        network
            qos network-policy-id
```

Use the following syntax to create a network QoS policy.

```
config>qos#
    network policy-id [network-policy-type network-policy-type]
        description description-string
    scope {exclusive|template}
        egress
            remarking
            remark <policy-id>
        ingress
            default-action fc {fc-name} profile {in|out}
            lsp-exp lsp-exp-value fc fc-name profile {in | out}
            fc {fc-name}
                meter {meter-id}
                multicast-meter {id}
            meter meter-id [multipoint]
                adaptation-rule cir {closest | max | min} pir {closest | max | min}
                cbs {size-in-kbits}
                mbs {size-in-kbits}
                mode {trtcm | srtcm}
                rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
            mpls-lsp-exp-profile policy-id
```

Output example

The following is a sample configuration output of a network **port** type QoS policy on the 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp.

```
*A:dut-d>config>qos>network# info detail
-----
    description "Default network-port QoS policy."
    scope template
    ingress
        default-action fc be profile out
        meter 1 create
        mode trtcm1
```

```
        adaptation-rule cir closest pir closest
        rate cir 0 pir max
        mbs default kbits
        cbs default kbits
    exit
    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 af41 fc h2 profile in
exit
egress
    no remarking
    fc af
        dscp-in-profile af11
        dscp-out-profile af12
        no dotlp
        dotlp-in-profile 3
        dotlp-out-profile 2
        no de-mark
    exit
    fc be
        dscp-in-profile be
        dscp-out-profile be
        no dotlp
        dotlp-in-profile 0
        dotlp-out-profile 0
        no de-mark
    exit
    fc ef
```

Output example

The following is a sample configuration output of a network **port** type QoS policy on the 7210 SAS-R6 and 7210 SAS-R12.

```
A:Dut-B>config>qos>network# info detail
-----
description "Default network-port QoS policy."
scope template
ingress
    default-action fc be profile out
    meter 1 create
        mode trtcml
        adaptation-rule cir closest pir closest
        rate cir 0 pir max
        mbs default kbits
        cbs default kbits
    exit
    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 af41 fc h2 profile in
exit
egress
    no remarking
```

```

        remark 1
        exit
-----
A:Dut-B>config>qos>network#

```

7.2.5 Default network policy values available in network mode

The default network policy for IP interfaces is identified as policy-id 2. Default policies cannot be modified or deleted. The following table lists default network policy parameters.

Table 47: Network policy defaults for policy ip-interface type

Field	Default
description	Default network QoS policy
scope	template
ingress	
default-action	fc be profile out (default action profile out is applicable only for port policies and not for ip-interface policies)
mpls-lsp-exp-profile	1
egress	
remarking	no
fc af:	
lsp-exp-in-profile	3
lsp-exp-out-profile	2
fc be:	
lsp-exp-in-profile	0
lsp-exp-out-profile	0
fc ef:	
lsp-exp-in-profile	5
lsp-exp-out-profile	5
fc h1:	
lsp-exp-in-profile	6
lsp-exp-out-profile	6
fc h2:	

Field	Default
lsp-exp-in-profile	4
lsp-exp-out-profile	4
fc h11:	
lsp-exp-in-profile	3
lsp-exp-out-profile	2
fc h12:	
lsp-exp-in-profile	1
lsp-exp-out-profile	1
fc nc:	
lsp-exp-in-profile	7
lsp-exp-out-profile	7

The following table lists default parameters for network QoS policy ip-interface type, LSP EXP-to-FC mapping on ingress. Color-aware policing is supported on network ingress.

Table 48: Default network QoS policy of ip-interface type, LSP EXP-to-FC mapping on Ingress

LSP EXP value	FC ingress	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

The default network policy for port is identified as *policy-id* 1. Default policies cannot be modified or deleted.

Output example

The following is a sample configuration output that shows the parameters for the default network **port** type QoS policy for the 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

```
*A:Dut-A>config>qos>network# info detail
-----
description "Default network-port QoS policy."
scope template
ingress
  default-action fc be profile out
  meter 1 create
    mode trtcm1
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs default kbits
    cbs default kbits
  exit
  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 af41 fc h2 profile in
exit
egress
  no remarking
  remark 1
exit
-----
*A:Dut-A>config>qos>network#
```

Output example

The following is a sample configuration output for the default remark policy used for dot1p and DSCP marking on the 7210 SAS-R6 and 7210 SAS-R12.

```
fc af
    dscp-in-profile af11
    dscp-out-profile af12
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dot1p-lsp-exp-in-profile
    no dot1p-lsp-exp-out-profile
    dot1p-in-profile 3
    dot1p-out-profile 2
exit
fc be
[no] de-mark
    [no] dot1p
    dscp-in-profile be
    dscp-out-profile be
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dot1p-lsp-exp-in-profile
    no dot1p-lsp-exp-out-profile
    dot1p-in-profile 0
    dot1p-out-profile 0
```

```
exit
fc ef
    dscp-in-profile ef
    dscp-out-profile ef
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dotlp-lsp-exp-in-profile
    no dotlp-lsp-exp-out-profile
    dotlp-in-profile 5
    dotlp-out-profile 5
exit
fc h1
    dscp-in-profile nc1
    dscp-out-profile nc1
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dotlp-lsp-exp-in-profile
    no dotlp-lsp-exp-out-profile
    dotlp-in-profile 6
    dotlp-out-profile 6
exit
fc h2
    dscp-in-profile af41
    dscp-out-profile af41
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dotlp-lsp-exp-in-profile
    no dotlp-lsp-exp-out-profile
    dotlp-in-profile 4
    dotlp-out-profile 4
exit
fc l1
    dscp-in-profile af21
    dscp-out-profile af22
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dotlp-lsp-exp-in-profile
    no dotlp-lsp-exp-out-profile
    dotlp-in-profile 3
    dotlp-out-profile 2
exit
fc l2
    dscp-in-profile cs1
    dscp-out-profile cs1
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dotlp-lsp-exp-in-profile
    no dotlp-lsp-exp-out-profile
    dotlp-in-profile 1
    dotlp-out-profile 1
exit
fc nc
    dscp-in-profile nc2
    dscp-out-profile nc2
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dotlp-lsp-exp-in-profile
    no dotlp-lsp-exp-out-profile
    dotlp-in-profile 7
    dotlp-out-profile 7
exit
```


7.2.6 Resource allocation for network QoS policy

This section describes the allocation of QoS resources for network QoS policies of the **ip-interface** and **port** type.

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

The resources are allocated to a network policy only when a port is configured for the IP interface. When a network port is configured, a default network QoS policy of the **port** type is applied.

For every FC in use, the system allocates two classification entries in hardware, provided the FC is configured to use both the unicast meter and multicast meter, or provided the default meter 9 is configured in the policy. If multiple match criteria entries map to the same FC, each of these are allocated two classification entries in hardware; for example, if there are two match-criteria entries that map to FC "af," a total of four classification entries are allocated in hardware, and if there are four match-criteria entries that map to FC "af," a total of eight classification entries are allocated in hardware.

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

The number of IP interfaces and network ports allowed is limited by the number of classification resources available in hardware, subject to the system limit on the number of IP interfaces and network or hybrid ports supported by the system.

Calculating the Number of QoS Resources

To calculate the number of QoS resources used by an IP interface, determine the following items:

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

Only the FCs used by the match-criteria classification entries are considered in the "number of FCs" and are therefore referred to as "FCs in use". In network policies of the **ip-interface** type, a default multipoint meter 9 is created in a policy; for policies of the **port** type, a default multipoint meter needs to be explicitly configured by the user, if required.

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

- If an FC is in use and is created without explicit meters, use default meter 1 for unicast traffic and default meter 9 (if configured) for all other traffic types (that is, broadcast, multicast, and unknown-unicast). This requires two classification entries in hardware. If default multipoint meter 9 is not configured, the FC uses the unicast meter for all traffic types. In this case, the FC requires a single classification entry in hardware.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter 9 (if configured) for all other traffic types. This requires two classification entries in hardware. If default multipoint meter 9 is not configured, the FC uses the unicast meter for all traffic types. In this case, the FC requires a single classification entry in hardware.
- If an FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and a multicast meter for all other kinds of traffic. This requires two classification entries in hardware.

Given the number of match criteria and the number of FCs in use, the following equation calculates the total number of classification entries (TC) per policy:

$$TC = \sum 2 * E(i)$$

i = nc, h1, ef, h2, l1, af, l2, be

where:

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



Note:

In the worst case, only 2 classification entries are used for each FC in a network policy, because only two types of traffic are supported.

Determining the number of policers or meters

Determine the number of policers or meters (TP) to use. A maximum of 16 meters per network policy are available.

The number of TPs used is the number of meters configured in the policy. Among that number, only those meters configured for use with an FC are considered during resource allocation. Meters that are created but not associated with an FC are not counted for resource allocation.

7.2.6.1 Network QoS policies resource usage examples



Note:

In the examples in this section, the profile configuration is not shown. In practice, users must configure the **mpls-lsp-exp-profile** policy and associate it with the network policy. Association of a profile policy with the network QoS policy does not change the resource calculation methodology shown in the following examples.

7.2.6.1.1 Example 1

Example

```
network 1 network-policy-type ip-interface create
  description "network-policy-1"
  ingress
    default-action fc be
    meter 1 create
    exit
    meter 9 multipoint create
    exit
  exit
  egress
    fc af
    exit
    fc be
    exit
    fc ef
    exit
    fc h1
    exit
    fc h2
    exit
    fc l1
```

```

        exit
        fc l2
    exit
    fc nc
    exit
exit

```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 0)nc + (2 * 0)h1 + (2 * 0)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 0)af + (2 * 0)l2 + (2 * 1)be = 2$$

The number of meters (TP) used is 2 (meters 1 and 9).

7.2.6.1.2 Example 2

Example

```

network 2 network-policy-type ip-interface create
description "network-policy-2"
    ingress
        default-action fc be
        meter 1 create
        exit
        meter 2 create
        exit
        meter 9 multipoint create
        exit
        meter 12 multipoint create
        exit
        fc "af" create
            meter 2
            multicast-meter 12
        exit
        lsp-exp 2 fc af
    exit
    egress
        fc af
        exit
        fc be
        exit
        fc ef
        exit
        fc h1
        exit
        fc h2
        exit
        fc l1
        exit
        fc l2
        exit
        fc nc
        exit
    exit
exit

```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 0)nc + (2 * 0)h1 + (2 * 0)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 1)af + (2 * 0)l2 + (2 * 1)be = 4$$

The number of meters (TP) user is 4 (meters 1, 2, 9, and 12)

7.2.6.1.3 Example 3

Example

```
network 3 network-policy-type ip-interface create
description "network-policy-3"
  ingress
    default-action fc be
    meter 1 create
    exit
    meter 2 create
    exit
    meter 9 multipoint create
    exit
    meter 12 multipoint create
    exit
    fc "af" create
      meter 2
      multicast-meter 12
    exit
    fc "be" create
      meter 2
      multicast-meter 12
    exit
    lsp-exp 2 fc af
  exit
  egress
    fc af
    exit
    fc be
    exit
    fc ef
    exit
    fc h1
    exit
    fc h2
    exit
    fc l1
    exit
    fc l2
    exit
    fc nc
    exit
  exit
exit
```

The number of classification entries (TC) used are calculated, as follows:

$$(2 * 0)nc + (2 * 0)h1 + (2 * 0)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 1)af + (2 * 0)l2 + (2 * 1)be = 4$$

The number of meters (TP) user is 2 (meters 2 and 12).

7.2.6.1.4 Example 4

Example

```
network 4 network-policy-type ip-interface create
description "network-policy-4"
  ingress
    default-action fc be
    meter 1 create
    exit
    meter 9 multipoint create
    exit
    lsp-exp 1 fc l2
    lsp-exp 2 fc af
    lsp-exp 3 fc af
    lsp-exp 4 fc h2
    lsp-exp 5 fc ef
    lsp-exp 6 fc h1
    lsp-exp 7 fc nc
  exit
  egress
    fc af
    exit
    fc be
    exit
    fc ef
    exit
    fc h1
    exit
    fc h2
    exit
    fc l1
    exit
    fc l2
    exit
    fc nc
    exit
  exit
exit
```

The number of Filter-Entries (TC) used is calculated, as follows:

$$(2 * 1)nc + (2 * 1)h1 + (2 * 1)ef + (2 * 1)h2 + (2 * 0)l1 + (2 * 2)af + (2 * 1)l2 + (2 * 1)be = 16$$

The number of meters (TP) used is 2 (meters 1 and 9).

7.2.6.1.5 Example 5

Example

```
network 5 network-policy-type ip-interface create
description "network-policy-5"
  ingress
    default-action fc be
    meter 1 create
    exit
    meter 2 create
    exit
```

```
meter 9 multipoint create
exit
meter 12 multipoint create
exit
fc "af" create
exit
fc "be" create
exit
fc "ef" create
exit
fc "h1" create
exit
fc "h2" create
exit
fc "l2" create
exit
fc "nc" create
exit
lsp-exp 1 fc l2
lsp-exp 2 fc af
lsp-exp 3 fc af
lsp-exp 4 fc h2
lsp-exp 5 fc ef
lsp-exp 6 fc h1
lsp-exp 7 fc nc
exit
egress
  fc af
  exit
  fc be
  exit
  fc ef
  exit
  fc h1
  exit
  fc h2
  exit
  fc l1
  exit
  fc l2
  exit
  fc nc
  exit
exit
```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 1)nc + (2 * 1)h1 + (2 * 1)ef + (2 * 1)h2 + (2 * 0)l1 + (2 * 2)af + (2 * 1)l2 + (2 * 1)be = 16$$

The number of meters (TP) used is 2 (meters 1 and 9). Note that meters 2 and 12 are not accounted for because they are not associated with any FC.

7.2.6.1.6 Example 6

Example

```
network 6 network-policy-type ip-interface create
  description "network-policy-6"

  ingress
```

```

        default-action fc be
        meter 1 create
        exit
        meter 2 create
        exit
        meter 3 create
        exit
        meter 9 multipoint create
        exit
        meter 12 multipoint create
        exit
        fc "af" create
            meter 2
            multicast-meter 12
        exit
        fc "be" create
        exit
        fc "ef" create
        exit
        fc "h1" create
            meter 3
        exit
        fc "h2" create
        exit
        fc "l2" create
        exit
        fc "nc" create
            meter 3
        exit
        lsp-exp 1 fc l2
        lsp-exp 2 fc af
        lsp-exp 3 fc af
        lsp-exp 4 fc h2
        lsp-exp 5 fc ef
        lsp-exp 6 fc h1
        lsp-exp 7 fc nc
    exit
    egress
        fc af
        exit
        fc be
        exit
        fc ef
        exit
        fc h1
        exit
        fc h2
        exit
        fc l1
        exit
        fc l2
        exit
        fc nc
        exit
    exit
exit

```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 1)nc + (2 * 1)h1 + (2 * 1)ef + (2 * 1)h2 + (2 * 0)l1 + (2 * 2)af + (2 * 1)l2 + (2 * 1)be = 16$$

The number of meters (TP) used is 5 (meters 1, 2, 3, 9, and 12).

7.2.6.1.7 Example 7

Example

```
network 2 network-policy-type ip-interface create
  description "Default network QoS policy."
  scope template
  ingress
    default-action fc be
    meter 1 create
      mode trtcm
      adaptation-rule cir closest pir closest
      rate cir 0 pir max
      mbs default
      cbs default
    exit
    meter 9 multipoint create
      mode trtcm
      adaptation-rule cir closest pir closest
      rate cir 0 pir max
      mbs default
      cbs default
    exit
    lsp-exp 0 fc be
    lsp-exp 1 fc l2
    lsp-exp 2 fc af
    lsp-exp 3 fc af
    lsp-exp 4 fc h2
    lsp-exp 5 fc ef
    lsp-exp 6 fc h1
    lsp-exp 7 fc nc
  exit
  egress
    no remarking
    fc af
      lsp-exp-in-profile 3
      lsp-exp-out-profile 2
    exit
    fc be
      lsp-exp-in-profile 0
      lsp-exp-out-profile 0
    exit
    fc ef
      lsp-exp-in-profile 5
      lsp-exp-out-profile 5
    exit
    fc h1
      lsp-exp-in-profile 6
      lsp-exp-out-profile 6
    exit
    fc h2
      lsp-exp-in-profile 4
      lsp-exp-out-profile 4
    exit
    fc l1
      lsp-exp-in-profile 3
      lsp-exp-out-profile 2
    exit
    fc l2
      lsp-exp-in-profile 1
      lsp-exp-out-profile 1
    exit
```



```
        fc nc
        lsp-exp-in-profile 7
        lsp-exp-out-profile 7
    exit
exit
```

The number of classification entries (TC) used is 2.

The number of meters (TP) used is 2.

7.2.6.1.8 Example 8

Example

```
network 8 network-policy-type ip-interface create
description "network-policy-8"
    ingress
        default-action fc nc
        meter 1 create
        exit
        meter 2 create
        exit
        meter 3 create
        exit
        meter 4 create
        exit
        meter 5 create
        exit
        meter 7 multipoint create
        exit
        meter 8 multipoint create
        exit
        meter 9 multipoint create
        exit
        meter 12 multipoint create
        exit
        fc "af" create
            meter 2
            multicast-meter 12
        exit
        fc "ef" create
            meter 4
            multicast-meter 8
        exit
        fc "h2" create
        exit
        fc "l2" create
            meter 3
            multicast-meter 7
        exit
        fc "nc" create
            meter 4
            multicast-meter 8
        exit
        lsp-exp 1 fc l2
        lsp-exp 3 fc af
        lsp-exp 5 fc ef
        lsp-exp 7 fc nc
    exit
egress
```

```

fc af
exit
fc be
exit
fc ef
exit
fc h1
exit
fc h2
exit
fc l1
exit
fc l2
exit
fc nc
exit
exit
exit

```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 2)nc + (2 * 0)h1 + (2 * 1)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 1)af + (2 * 1)l2 + (0 * 0)be = 10$$

The numbers of meters (TP) used is 6 (meters 2, 3, 4, 7, 8, and 12).

7.3 Network QoS policy in access-uplink mode

The network QoS policy consists of an ingress and egress component. For 7210 SAS-T devices operating in access-uplink mode, network policy is available for use. The ingress component of the policy defines how dot1p bits are mapped to internal forwarding class and profile state (DSCP is not available for use). The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the system.

The mapping on each access uplink port defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the access uplink ports. It also defines the rate-limiting parameters for the traffic mapped to each forwarding classes. Traffic mapped to each forwarding class can be rate limited using separate meters for each unicast and multipoint traffic.

The egress component of the network QoS policy provides an option to define either dot1p bits and IP DSCP marking values or both associated with each forwarding class. By default, network QoS policy remarking is always disabled. If the egressing packet originated on an ingress SAP, the egress QoS policy also defines the dot1p bit marking based on the forwarding class and the profile state. The default map of FC-dot1p marking is as shown in default network QoS policy of type "port", where in the "policy-id" is equal to 1. All non-default network QoS policies inherits the FC-dot1p map.

Network policy-id 1 exists as the default policy and is applied to access uplink ports. The network policy-id 1 cannot be modified or deleted. It defines the default dot1p-to-FC mapping and default meters for unicast and multipoint meters 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. The **no** form of this command reverts the object to the default values.

Changes made to a policy are applied immediately to all ports where the policy is applied. For this reason, when a policy requires several changes, it is recommended that you copy the policy to a work area policy-

id. The work-in-progress copy can be modified until all the changes are made and then the original policy-id can be overwritten with the **config qos copy** command.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your devices, see the "CLI Usage" chapter in the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide*.

7.3.1 Basic configurations

A basic network QoS policy must conform to the following:

- Each network QoS policy must have a unique policy ID.
- The network must specify the default-action.
- The network must have a QoS policy scope of **template** or **exclusive**.

7.3.2 Configuring network policy for access-uplink mode

To create an network QoS policy when operating in access-uplink mode, define the following:

- A network policy ID value. The system does not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- You can modify egress dot1p and/or IP DSCP marking map. 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 dot1p bit and/or IP DSCP value mapping.
 - **forwarding class criteria**
The forwarding class name represents an egress queue. Specify forwarding class criteria to define the marking criteria of packets flowing through it.
 - **DSCP and dot1p**
The DSCP and dot1p value to use for all packets requiring marking that egress on this forwarding class queue that are in or out of profile.
- **Ingress criteria**
Specifies dot1p to forwarding class mapping for all packets.
 - **default action**
Defines the default action to be taken for packets that have an undefined DSCP or dot1p bits set. The default-action specifies the forwarding class to which such packets are assigned.
 - **DSCP and dot1p**
Creates a mapping between the DSCP and dot1p bits of the access uplink port ingress traffic and the forwarding class. Ingress traffic that matches the specified DSCP and dot1p bits will be assigned to the corresponding forwarding class.

The following commands associated a network QoS policy with the access-uplink port.

```
config>port  
ethernet
```

```
access
  uplink
    qos network-policy-id
```

Use the following syntax to create a network QoS policy for 7210 SAS-T in access uplink mode.

Example

```
A:MTU>config>qos>network# info detail
-----
description "Default network-port QoS policy."
scope template
ingress
  default-action fc be profile out
  meter 1 create
    mode trtcml
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs default kbits
    cbs default kbits
  exit
  meter 9 multipoint create
    mode trtcml
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs default kbits
    cbs default kbits
  exit
  dot1p 0 fc be profile out
  dot1p 1 fc l2 profile in
  dot1p 2 fc af profile out
  dot1p 3 fc af profile in
  dot1p 4 fc h2 profile in
  dot1p 5 fc ef profile in
  dot1p 6 fc h1 profile in
  dot1p 7 fc nc profile in
exit
egress
  no remarking
  fc af
    dscp-in-profile af11
    dscp-out-profile af12
    no dot1p
    dot1p-in-profile 3
    dot1p-out-profile 2
    no de-mark
  exit
  fc be
    dscp-in-profile be
    dscp-out-profile be
    no dot1p
    dot1p-in-profile 0
    dot1p-out-profile 0
    no de-mark
  exit
  fc ef
    dscp-in-profile ef
    dscp-out-profile ef
    no dot1p
    dot1p-in-profile 5
A:MTU>config>qos>network#
```

7.3.3 Default network policy values available in access-uplink mode

Output example

The following sample output is of a default network port policy on 7210 SAS-T in access-uplink mode.

```
*7210SAST>config>qos>network# info detail
-----
description "Default network-port QoS policy."
scope template
ingress
  default-action fc be profile out
  meter 1 create
    mode trtcml
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs default kbits
    cbs default kbits
  exit
  meter 9 multipoint create
    mode trtcml
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs default kbits
    cbs default kbits
  exit
  dot1p 0 fc be profile out
  dot1p 1 fc l2 profile in
  dot1p 2 fc af profile out
  dot1p 3 fc af profile in
  dot1p 4 fc h2 profile in
  dot1p 5 fc ef profile in
  dot1p 6 fc h1 profile in
  dot1p 7 fc nc profile in
exit
egress
  no remarking
  fc af
    dscp-in-profile af11
    dscp-out-profile af12
    no dot1p
    dot1p-in-profile 3
    dot1p-out-profile 2
    no de-mark
  exit
  fc be
    dscp-in-profile be
    dscp-out-profile be
    no dot1p
    dot1p-in-profile 0
    dot1p-out-profile 0
    no de-mark
  exit
  fc ef
    dscp-in-profile ef
    dscp-out-profile ef
    no dot1p
    dot1p-in-profile 5
    dot1p-out-profile 5
    no de-mark
  exit
  fc h1
```

```

        dscp-in-profile nc1
        dscp-out-profile nc1
        no dot1p
        dot1p-in-profile 6
        dot1p-out-profile 6
        no de-mark
    exit
fc h2
        dscp-in-profile af41
        dscp-out-profile af41
        no dot1p
        dot1p-in-profile 4
        dot1p-out-profile 4
        no de-mark
    exit
fc l1
        dscp-in-profile af21
        dscp-out-profile af22
        no dot1p
        dot1p-in-profile 3
        dot1p-out-profile 2
        no de-mark
    exit
fc l2
        dscp-in-profile cs1
        dscp-out-profile cs1
        no dot1p
        dot1p-in-profile 1
        dot1p-out-profile 1
        no de-mark
    exit
fc nc
        dscp-in-profile nc2
        dscp-out-profile nc2
        no dot1p
        dot1p-in-profile 7
        dot1p-out-profile 7
        no de-mark
    exit
exit
-----

```

7.4 DSCP and dot1p marking for CPU self-generated traffic

DSCP and dot1p marking for CPU self-generated traffic is not user-configurable, except on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 (see [QoS for self-generated \(CPU\) traffic on network interfaces for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#)). [Table 49: Default CPU QoS values for DSCP and dot1p marking](#) lists the default CPU QoS values for DSCP and dot1p marking.



Note:

- Protocols such as BGP, RSVP, TLDP, OSPF, and IS-IS are not supported on 7210 SAS platforms operating in access-uplink mode.
- For PTP messages, based on the message type (event or non-event), the DSCP value used is either 0x30 (h1) or 0x38 (nc), and the dot1p value is always 7.

- DSCP and dot1p values in the table are applicable when remarking is disabled at the port level.

Table 49: Default CPU QoS values for DSCP and dot1p marking

Protocol	IPv4	DSCP marking	Dot1p marking	Default FC	DSCP values (decimal)	Dot1p values
ARP	N/A	N/A	Yes	NC	—	7
BGP	Yes	Yes	Yes	NC	48	7
Cflowd	Yes	Yes	Yes	NC	48	7
CFM	N/A	N/A	Yes	NC	—	7
DNS	Yes	Yes	Yes	H2	34	4
FTP	Yes	Yes	Yes	H2	34	4
ICMP Req	Yes	Yes	Yes	NC	0	7
ICMP Res	Yes	Yes	Yes	NC	0	7
ICMP Unreach	Yes	Yes	Yes	NC	0	7
IGMP	Yes	Yes	Yes	NC	48	7
IS-IS	Yes	No	Yes	NC	—	7
MSDP	Yes	Yes	Yes	NC	48	7
NTP	Yes	Yes	Yes	NC	48	7
OSPF	Yes	Yes	Yes	NC	48	7
PIM (SSM)	Yes	Yes	Yes	NC	48	7
PTP	Yes	Yes	Yes	H2	48	7
RADIUS	Yes	Yes	Yes	H2	34	4
RSVP	Yes	Yes	Yes	NC	48	7
SCP	Yes	Yes	Yes	H2	34	4
SNMP	Yes	Yes	Yes	H2	34	4
SSH	Yes	Yes	Yes	H2	34	4
STP	N/A	N/A	Yes	NC	—	7
SYSLOG	Yes	Yes	Yes	H2	34	4
TACACS	Yes	Yes	Yes	H2	34	4

Protocol	IPv4	DSCP marking	Dot1p marking	Default FC	DSCP values (decimal)	Dot1p values
TACPLUS	Yes	Yes	Yes	H2	34	4
Telnet	Yes	Yes	Yes	H2	34	4
TFTP	Yes	Yes	Yes	H2	34	4
TLDP	Yes	Yes	Yes	NC	48	7
Trace route	Yes	Yes	Yes	NC	0	7

7.4.1 QoS for self-generated (CPU) traffic on network interfaces for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Differentiated services code point (DSCP), forwarding class (FC), and IEEE 802.1p values can be specified for use by protocol packets generated by the node. This enables prioritization or deprioritization of supported protocols (as required).

DSCP marking for internally generated control and management traffic should be used for a specified 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 and IS-IS 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 and an FC, the 802.1p and MPLS EXP bits are 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 the packet is egressing.

Configuring self-generated QoS is supported in the base router and VPRN service contexts.

The default values for self-generated traffic on network interfaces are the following:

- routing protocols (for example, OSPF and BGP)
 - FC - Network Control (NC)
 - DSCP value - NC1 (not applicable for ARP and IS-IS)
 - 802.1p value - according to the egress QoS policy (7 by default)
- management protocols (for example, SSH and SNMP)
 - FC - Network Control (NC)
 - DSCP value - AF41
 - 802.1p value - according to the egress QoS policy (7 by default)



Note:

- ICMP echo requests (type 8) and ICMPv6 echo requests (type 128) initiated from the router use the DSCP value that is specified in the **sgt-qos** command. The FC value is NC by default, or the value that is specified in the **ping** command parameter **fc fc-name**.
- Configurable values for BFD are not supported.

- On access SAP egress and access port egress, when remarking is not enabled, the dot1p value for all IP packets generated by the node is set to zero. To enable dot1p marking, remarking must be enabled.

7.4.2 Default DSCP mapping table

The following table lists the DSCP mapping between DSCP name and DSCP values (decimal, hexadecimal, and binary) and label.

Table 50: 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 15	0			

¹⁵ The default FC mapping is used for all DSCP names/values for which there is no explicit FC mapping.

7.5 Service management tasks

This section provides information about service management tasks.

7.5.1 Deleting QoS policies

A network policy is associated by default with IP interfaces and network ports for devices operating in network mode. A network policy is associated by default with access uplink ports for devices in access uplink mode.

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

7.5.2 Remove a policy from the QoS configuration

Use the following syntax to delete a network policy.

```
config>qos# no network network-policy-id
```

7.5.3 Copying and overwriting network policies

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

```
config>qos# copy network source-policy-id dest-policy-id [overwrite]
```

Output example

The following is a sample configuration output.

```
A:ALA-12>config>qos# info detail
-----
...
    network 1 create
      description "Default network QoS policy."
      scope template
      ingress
      default-action fc be profile out
...
    network 600 create
      description "Default network QoS policy."
      scope template
      ingress
      default-action fc be profile out
...
    network 700 create
      description "Default network QoS policy."
      scope template
      ingress
      default-action fc be profile out
```

```
...
-----
A:ALA-12>config>qos#
```

7.5.4 Editing QoS policies

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

7.6 Network QoS policy command reference

7.6.1 Command hierarchies

- Configuration commands for MPLS EXP profile map (7210 SAS platforms operating in network mode)
- Configuration commands (7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE configured in network mode)
- Configuration commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)
- Configuration commands (access-uplink mode)
- Self-generated traffic commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)
- Operational commands (network mode or access-uplink mode)
- Show commands (network mode or access-uplink mode)
- Show commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)

7.6.1.1 Configuration commands for MPLS EXP profile map (7210 SAS platforms operating in network mode)

```
config
- qos
  - [no] mpls-lsp-exp-profile-map policy-id [create]
    - description description-string
    - no description
    - lsp-exp lsp-exp-value profile {in | out}
    - no lsp-exp
  - [no] use-global-mpls-lsp-exp-profile policy-id
```

7.6.1.2 Configuration commands (7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE configured in network mode)

```
config
- qos
```

```
- [no] network network-policy-id [create]
- [no] network network-policy-id [create] [network-policy-type {ip-interface | port}]
  - description description-string
  - no description
  - scope {exclusive | template}
  - no scope
  - egress
    - no remark
    - remark policy-id
    - remarking
    - no remarking
  - ingress
    - default-action fc fc-name profile {in | out | use-dei}
    - dotlp dotlp-priority fc fc-name profile {in | out | use-dei}
    - no dotlp dotlp-priority
    - [no] fc fc-name [create]
      - meter meter-id
      - no meter
      - multicast-meter meter-id
      - no multicast-meter
    - dscp dscp-name fc fc-name profile {in | out}
    - no dscp dscp-name
    - lsp-exp lsp-exp-value fc fc-name
    - no lsp-exp lsp-exp-value
    - meter meter-id [multipoint] [create]
    - no meter meter-id
      - adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
      - no adaptation-rule
      - cbs size [kbits | bytes | kbytes]
      - no cbs
      - mbs size [kbits | bytes | kbytes]
      - no mbs
      - mode mode
      - no mode
      - rate cir-rate-in-kbps [pir pir-rate-in-kbps]
      - no rate
    - [no] mpls-lsp-exp-profile policy-id
```

7.6.1.3 Configuration commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)

```
config
- qos
  - [no] network network-policy-id [create] [network-policy-type {ip-interface | port}]
    - description description-string
    - no description
    - egress
      - remark policy-id
      - no remark
      - remarking
      - no remarking
    - ingress
      - default-action fc fc-name profile {in | out | use-dei}
      - dotlp dotlp-priority fc fc-name profile {in | out | use-dei}
      - no dotlp dotlp-priority
      - [no] fc fc-name [create]
        - meter meter-id
        - no meter
        - multicast-meter meter-id
        - no multicast-meter
      - dscp dscp-name fc fc-name profile {in | out}
      - no dscp dscp-name
```

```
- lsp-exp lsp-exp-value fc fc-name
- no lsp-exp lsp-exp-value
- mpls-lsp-exp-profile policy-id
- no mpls-lsp-exp-profile
- meter meter-id [multipoint] [create]
- no meter meter-id
  - adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
  - no adaptation-rule
  - cbs size [kbits | bytes | kbytes]
  - no cbs
  - mbs size [kbits | bytes | kbytes]
  - no mbs
  - mode mode
  - no mode
  - rate cir-rate-in-kbps [pir pir-rate-in-kbps]
  - no rate
- scope {exclusive | template}
- no scope
```

7.6.1.4 Configuration commands (access-uplink mode)

```
config
- qos
  - [no] network network-policy-id [network-policy-type] {ip-interface | port} (only type
  port is supported)
    - description description-string
    - no description
    - no scope {exclusive | template}
    - egress
      - [no] fc fc-name
        - [no] de-mark [force de-value]
        - dotlp dotlp-priority
        - no dotlp
        - dscp-in-profile dscp-name
        - no dscp-in-profile
        - dscp-out-profile dscp-name
        - no dscp-out-profile
        - dotlp-in-profile dotlp-priority
        - no dotlp-in-profile
        - dotlp-out-profile dotlp-priority
        - no dotlp-out-profile
      - [no] remarking {use-dotlp | use-dscp | all}
    - ingress
      - default-action fc fc-name profile {in | out | use-dei}
      - dotlp dotlp-priority fc fc-name profile {in | out | use-dei}
      - no dotlp dotlp-priority
      - [no] fc fc-name [create]
        - meter meter-id
        - no meter
        - multicast-meter meter-id
        - no multicast-meter
      - meter meter-id [multipoint] [create]
      - no meter meter-id
        - adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
        - no adaptation-rule
        - cbs size-in-kbits
        - no cbs
        - mbs size-in-kbits
        - no mbs
        - mode {trtcm1 | trtcm2 | srtcm}
        - no mode
```

```
- rate cir-rate-in-kbps [pir pir-rate-in-kbps]
- no rate
```

7.6.1.5 Self-generated traffic commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)

```
config
- router
  - sgt-qos
    - application dscp-app-name dscp {dscp-value | dscp-name}
    - application dot1p-app-name dot1p dot1p-priority
    - no application {dscp-app-name | dot1p-app-name}
    - dscp dscp-name fc fc-name
    - no dscp dscp-name
- service
  - vprn
    - sgt-qos (See Note below)
      - 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
```



Note:

For descriptions of the **config service vprn sgt-qos** commands, refer to the VPRN Service Configuration Commands section in the *7210 SAS-Mxp, S, Sx, T Services Guide* and the *7210 SAS-R6, R12 Services Guide*.

7.6.1.6 Operational commands (network mode or access-uplink mode)

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

7.6.1.7 Show commands (network mode or access-uplink mode)

```
show
- qos
  - network policy-id [detail]
  - network [network-policy-id] association
  - network [network-policy-id] [detail]
  - mpls-lsp-exp-profile-map [policy-id] [detail] (available only in Network Mode)
```

7.6.1.8 Show commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)

```
show
- qos
  - dscp-table [value dscp-value]
show
- router [router-instance]
- router service-name service-name
  - sgt-qos
```

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

7.6.2 Command descriptions

- [Configuration commands](#)
- [Self-generated traffic commands \(7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12\)](#)
- [Operational commands](#)
- [Show commands](#)

7.6.2.1 Configuration commands

- [Generic commands](#)
- [Network QoS policy commands](#)
- [Network QoS policy commands \(7210 SAS-T in access-uplink mode\)](#)
- [Network ingress QoS policy commands](#)
- [Network egress QoS policy commands](#)

7.6.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>network

config>qos>mpls-lsp-exp-profile-map

Platforms

Supported on all 7210 SAS platforms as described in this document, including platforms configured in the access-uplink operating mode

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

description-string

Specifies 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.

7.6.2.2 Network QoS policy commands

network

Syntax

network *network-policy-id* [**create**] [**network-policy-type** {**ip-interface** | **port**}]

no network *network-policy-id*

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates or edits a QoS network policy. The network policy defines the treatment packets receive as they ingress and egress the network port and network IP interface in network mode of operation.

The QoS network policy consists of an ingress and egress component. The ingress component of the policy defines how packet header priority 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 7210 SAS. 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. It also defines the rate-limiting parameters for the traffic mapped to each forwarding classes. Traffic mapped to each forwarding class can be rate limited using separate meters for unicast and multipoint traffic.

The egress component of the network QoS policy defines forwarding class and profile state to packet header priority bit values for traffic to be transmitted into the core network. If the egressing packet originated on an ingress SAP, the parameter is always enabled for the network port. The egress QoS policy also defines the dot1p bit marking based on the forwarding class and the profile state.

In network mode, network policy ID 2 exists as the default policy that is applied to all IP interface by default. The network policy ID 2 cannot be modified or deleted. It defines the default LSP EXP-to-FC mapping and default meters for unicast traffic and optional multipoint meters for BUM traffic on the ingress MPLS packets. For the egress, it defines eight forwarding classes which define LSP EXP values and the packet marking behavior.

In network mode, Network *policy-id* 1 exists as the default policy that is applied to all network ports by default. This default policy cannot be modified or deleted. It defined the default DSCP-to-FC mapping and

default unicast meters for ingress IP traffic. For the egress, it defines the forwarding class to dot1p and DSCP values and the packet marking criteria.

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

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

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

Default

system default network policy 1

Parameters

network-policy-id

Specifies the policy on the interface or port.

Values 1 to 65535

network-policy-type

Specifies the type of the policy. This parameter defines where this network policy can be applied.

Values **ip-interface** — Specifies only EXP-based classification rules and marking values. It can only be associated with an IP interface. It can be used only when the device is operating in network mode.
port — Specifies only DSCP and dot1p classification rules and marking values. It can only be associated with a network port or hybrid port.

mpls-lsp-exp-profile-map

Syntax

mpls-lsp-exp-profile-map *policy-id* [**create**]

no mpls-lsp-exp-profile-map

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a new **mpls-lsp-exp-profile-map** policy. The policy specifies the profile to assign to the packet based on the MPLS LSP EXP bits value matched in the MPLS packet received on a network IP interface.

The assigned profile is available for use by the meter or policer associated with FC in the network policy attached to this IP interface.

The policy is associated with network policy attached to a network IP interface.

Default

1 (default mpls-lsp-exp-profile-map policy "1")

Parameters

policy-id

Specifies the policy ID on the 7210 SAS.

Values 1 to 65535

create

Keyword to create a policy.

lsp-exp

Syntax

lsp-exp *lsp-exp-value*

no **lsp-exp**

Context

config>qos> mpls-lsp-exp-profile-map

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

Ingress traffic that matches the specified LSP EXP bits will be assigned the corresponding profile.

Multiple commands can be entered to define the association of some or all eight LSP EXP bit values to the profile. For undefined values, packets are assigned the profile value **out**.

The **no** form of this command removes the association of the LSP EXP bit value to the profile value. The default profile value **out** then applies to that LSP EXP bit pattern.

Parameters

lsp-exp-value

Specifies a 3-bit LSP EXP bit value, expressed as a decimal integer.

Values 0 to 7

use-global-mpls-lsp-exp-profile

Syntax

use-global-mpls-lsp-exp-profile *policy-id*

no use-global-mpls-lsp-exp-profile

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command associates the **mpls-lsp-exp-profile-map** policy for use with LDP LSPs. When color-aware metering is in use for the IP interface, the policy specified here provides the profile to assign to the MPLS packets received on any of the network IP interface in use in the system. The MPLS EXP bits in the received packet are matched for assigning the profile.

On system boot-up, the policy ID is set to the default policy ID 1. The user can modify it to use the policy of their choice.

For LDP LSP traffic and segment routing (SR), the system always uses the global **mpls-lsp-exp-profile-map** policy. For RSVP LSP traffic, system uses the **mpls-lsp-exp-profile-map** policy associated with the network policy. For consistent QoS treatment, Nokia highly recommends using a single **mpls-lsp-exp-profile-map** policy for all network policies when the FRR facility or a mix of LDP, SR, and RSVP is in use, or when the PHP is enabled.

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

Default

1

Parameters

policy-id

Specifies the **mpls-lsp-exp-profile-map** policy to use.

Values 1 to 65535

mpls-lsp-exp-profile

Syntax

mpls-lsp-exp-profile *policy-id* [**create**]

no mpls-lsp-exp-profile

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the **mpls-lsp-exp-profile-map** policy to use for assigning profile values for packets received on this IP interface.



Note:

For LDP LSP traffic, the system uses the global **mpls-lsp-exp-profile-map** policy. For RSVP LSP traffic, the system uses the **mpls-lsp-exp-profile-map** policy that is associated with the network policy. For consistent QoS treatment, Nokia highly recommends using a single **mpls-lsp-exp-profile-map** policy for all the network policies when FRR facility is in use.

The **no** form of this command removes the policy.

Parameters

policy-id

Specifies the policy.

Values 1 to 65535

create

Keyword to create a policy.

7.6.2.3 Network QoS policy commands (7210 SAS-T in access-uplink mode)

network

Syntax

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

Context

config>qos

Platforms

7210 SAS-T (access-uplink mode)

Description

This command creates or edits a QoS network policy. The network policy defines the treatment packets receive as they ingress and egress the access-uplink port. Only network-policy-type **port** is supported in access-uplink mode.

The QoS network policy consists of an ingress and egress component. The ingress component of the policy defines how dot1p 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 7210 SAS. The mapping on each access-uplink port defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the port. It also defines the rate-limiting parameters for the traffic mapped to each forwarding classes. Traffic mapped to each forwarding class can be rate limited using separate meters for each uni-cast and multipoint traffic.

The egress component of the network QoS policy defines the queuing parameters associated with each forwarding class. There are eight queues per port on the egress. Each of the forwarding classes is associated with a queue on each access-uplink port. 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, the parameter is always enabled for the access-uplink port, the egress QoS policy also defines the dot1p bit marking based on the forwarding class and the profile state.

The network policy ID 1 cannot be modified or deleted. It defines the default dot1p-to-FC mapping and default meters for unicast traffic and optional multipoint meters for the ingress. For the egress, it defines eight forwarding classes which represent individual queues and the packet marking criteria.

If a new network policy is created (for instance, policy ID 2), only the default action and default meters for unicast and multipoint traffic and egress forwarding class parameters are identical to the default policy. A new network policy does not contain the default dot1p-to-FC mapping for network QoS policy of type port. The default network policy can be copied using the **copy** command to create a new network policy that includes the default ingress dot1p or DSCP to FC mapping (as appropriate).

You can modify parameters or use the **no** modifier to remove an object from the configuration.

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

The **no** form of this command deletes the network policy. A policy cannot be deleted until it is removed from all entities where it is applied. The default network policy ID1 cannot be deleted.

Default

System Default Network Policy 1

Parameters

network-policy-id

Specifies the policy on the 7210 SAS.

Values 1 to 65535

network-policy-type

Specifies the type of network policy. This parameter defines where this network policy can be applied. In the access-uplink mode, only **port** type must be used and attached to an access-uplink port.

Values **ip-interface** — Specifies only EXP-based classification rules and marking values. It can only be associated with an IP interface. It can be used only when the device is operating in network mode.

port — Specifies only DSCP and dot1p classification rules and marking values. It can only be associated with a network port or hybrid port.

7.6.2.4 Network ingress QoS policy commands

fc

Syntax

fc *fc-name* [create]

no fc *fc-name*

Context

config>qos>network>ingress

config>qos>network>egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description



Note:

The **config>qos>network>egress** context is only supported on the 7210 SAS-T (access-uplink mode).

This command creates a class instance of the forwarding class. After the *fc-name* is created, classification actions can be applied and it can be used in match classification criteria.

The **no** form of this command removes all the explicit meter mappings for *fc-name* forwarding types. The meter mappings revert to the default meters for *fc-name*.

Default

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

Parameters

fc-name

Specifies a case-sensitive, system-defined forwarding class name for which policy entries will be created.

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

create

Creates the forwarding class. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

ingress

Syntax

ingress

Context

config>qos>network policy-id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in the access-uplink mode

Description

This command creates or edits policy entries that specify the *lsp-exp-value* to forwarding class mapping for all MPLS packets.

When pre-marked packets ingress on a network port, the QoS treatment through the 7210 SAS-based on the mapping defined under the current node.

default-action

Syntax

default-action fc fc-name [profile {in | out | use-dei}]

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in the access-uplink mode

Description

This command defines or edits the default action to be taken for packets that have an undefined LSP EXP (only when devices is operating in network mode) or dot1p bits (for 7210 SAS-T in access-uplink mode) bits set. The **default-action** command specifies the forwarding class to which such packets are assigned.

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

Default

default-action fc be profile out

Parameters

fc *fc-name*

Specifies the forwarding class name. All packets with LSP EXP (only when devices is operating in network mode) or dot1p bits (for 7210 SAS-T in access-uplink mode) that are not defined will be placed in this forwarding class.

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

profile {in | out | use-dei}

Specifies that all packets assigned to this forwarding class are considered in or out of profile based on this command. A value must be specified when the **profile** keyword is used in the command. If the profile is not assigned to a forwarding class, the packets of that FC are treated as out-of-profile packets.

Values **in** — Defines the packet profile as in-profile.
out — Defines the packet profile as out-of-profile
use-dei — Specifies that DEI is used to determine the initial profile of the packet

dot1p

Syntax

dot1p *dot1p-priority* **fc** *fc-name* **profile** {in | out | use-dei}

no dot1p

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in the access-uplink mode

Description

This command explicitly sets the forwarding class or enqueueing priority and profile of the packet when a packet has the *dot1p-priority* specified. Adding a dot1p rule on the policy forces packets that match the *dot1p-priority* specified to be assigned to the forwarding class and enqueueing priority and profile of the packet based on the parameters included in the dot1p rule.

The *dot1p-priority* is derived from the most significant three bits in the IEEE 802.1 q or IEEE 802.1p header. The three dot1p bits define eight Class-of-Service (CoS) values commonly used to map packets to per-hop 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 from all ingress SAP ports using the policy.

Parameters

dot1p-priority

Specifies the unique IEEE 802.1p value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-priority* value, the previous forwarding class is completely overridden by the new parameters.

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

Values 0 to 7

fc fc-name

Specifies a 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-dei}

Specifies that all packets that are assigned to this forwarding class will be considered in-profile or out-of-profile based on this command or to use the default. In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets.

Values **in** — Defines the packet profile as in-profile.

out — Defines the packet profile as out-of-profile.

use-dei — Specifies that DEI is used to determine the initial profile of the packet

meter

Syntax

meter *meter-id* [**multipoint**] [**create**]

no meter *meter-id*

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in the access-uplink mode

Description

This command configures an ingress Network QoS policy meter. The **meter** command allows the creation of multipoint meters. Only multipoint meters can receive ingress packets that need to be sent to multiple destinations

Multipoint meters are for traffic bound to multiple destinations. Within non-multipoint services, such as Epipe services, all traffic is considered unicast because of the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service will not be mapped to a multipoint service meter.

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

Default

meter 1 (for unicast traffic)

meter 9 multipoint (for all other traffic, other than unicast traffic)

Parameters

meter-id

Specifies the meter within the policy. The value 9 is reserved for the default multipoint meter.

Values 1 to 16 (For network policy of type **ip-interface**)
 1 to 8 (For network policy of type **port**)

multipoint

Specifies that this *meter-id* is for multipoint forwarded traffic only. This *meter-id* can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. If you attempt to map forwarding class unicast traffic to a multipoint queue, an error is generated and no changes are made to the current unicast traffic queue mapping.

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

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

meter

Syntax

meter *meter-id*

no meter

Context

config>qos>network>ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in the access-uplink mode

Description

This command overrides the default unicast forwarding type meter mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy as a non-multipoint meter before the mapping can be made. After the forwarding class mapping is executed, all unicast traffic on a port using this policy is forwarded using the *meter-id*.

The **no** form of this command sets the unicast (point-to-point) *meter-id* back to the default meter for the forwarding class (meter 1).

Default

meter 1

Parameters

meter-id

Specifies the meter. The specified parameter must be an existing, non-multipoint meter defined in the **config>qos>network>ingress** context.

Values 1 to 16

multicast-meter

Syntax

multicast-meter *meter-id*

no multicast-meter

Context

config>qos>network>ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

This command can be used with a network policy of type **ip-interface**, and on a network port when multicast is enabled.

The **no** form of this command reverts the multicast forwarding type *meter-id* to the default meter for the forwarding class.

Default

9

Parameters

meter-id

Specifies the multicast meter. The specified parameter must be an existing, multipoint meter defined in the **config>qos>network>ingress** context.

Values 1 to 16

dscp

Syntax

dscp *dscp-name* **fc** *fc-name* **profile** {**in** | **out**}

no dscp

Context

config>qos>network policy-id>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

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 DSCP-to-FC association. The **default-action** then applies to that code point value.

Parameters

dscp-name

Specifies 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.

[Table 51: Default DSCP names to DSCP value mapping table](#) shows the available system-defined names. The system-defined names must be referenced as all lowercase exactly as listed in the first column in [Table 51: Default DSCP names to DSCP value mapping table](#) and [Table 52: Default class selector code points to DSCP value mapping table](#) below.

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

The actual mapping is being done on the *dscp-value*, not the *dscp-name* that references the *dscp-value*. If a second *dscp-name* that references the same *dscp-value* is mapped

within the policy, an error will occur. The second name will not be accepted until the first name is removed.

Table 51: Default DSCP names to DSCP value mapping table

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

Table 52: Default class selector code points to DSCP value mapping table

DSCP name	DSCP value decimal	DSCP value hexadecimal	DSCP value binary
cs7	56	0x38	0b111000
cs6	48	0x30	0b110000
cs5	40	0x28	0b101000
cs4	32	0x20	0b100000
cs3	24	0x18	0b011000

DSCP name	DSCP value decimal	DSCP value hexadecimal	DSCP value binary
cs2	16	0x10	0b010000
cs1	08	0x8	0b001000

fc *fc-name*

Specifies the *fc-name* with which the code point will be associated.

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

profile {in | out}

Keyword 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.

**Note:**

- 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** — Defines the packet profile as in-profile.
 out — Defines the packet profile to be out-of-profile.

lsp-exp

Syntax

lsp-exp *lsp-exp-value* **fc** *fc-name*

no lsp-exp *lsp-exp-value*

Context

config>qos>network policy-id>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

Parameters

lsp-exp-value

Specifies the LSP EXP values to be associated with the forwarding class.

Values 0 to 7 (decimal representation of three EXP bit field)

fc *fc-name*

Specifies the *fc-name* that the EXP bit pattern will be associated with.

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

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]

no adaptation-rule

Context

config>qos>network>ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in the access-uplink mode

Description

This command defines the method used by the system to derive the operational CIRs and PIRs when the meter is provisioned in hardware. For the **cir** and **pir** parameters, the system attempts to find the best operational rate depending on the defined constraint.



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

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. When a specific **adaptation-rule** is removed, the default constraints for **cir** and **pir** apply.

Default

adaptation-rule cir closest pir closest

Parameters

cir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced to adapt the CIR defined using the **meter meter-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used to derive the operational CIR for the meter. When the **cir** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive.

See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) and [Table 34: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Sx 10/100GE](#) for more information about supported hardware step-size rates.

Default	closest
Values	<p>max — Specifies that the operational CIR value is equal to or less than the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.</p> <p>min — Specifies that the operational CIR value is equal to or greater than the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.</p> <p>closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.</p>

pir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced to adapt the PIR defined using the **meter meter-id rate** command. The **pir** parameter requires a qualifier that defines the constraint used to derive the operational PIR for the meter. When the **rate** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive.

See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) and [Table 34: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Sx 10/100GE](#) for information about supported hardware step-size rates.

Default	closest
Values	<p>max — Specifies that the operational PIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.</p> <p>min — Specifies that the operational PIR value is equal to or greater than the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.</p> <p>closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.</p>

cbs

Syntax

cbs *size* [**kbits** | **bytes** | **kbytes**]

no cbs

Context

config>qos>network>ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command overrides the default CBS for the meter. The committed burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the CBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter-configured parameters.



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

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

Default

32 kbits

Parameters

size

Specifies the number of kilobits or kilobytes or bytes reserved for the meter. For example, if a value of 100 kbits is desired, then enter the value 100. The bucket size is rounded off to the next highest 4096 bytes boundary.

Values	kbits — 4 to 2146959, default
	bytes — 512 to 274810752
	kbytes — 1 to 268369

[kbits | bytes | kbytes]

Specifies the unit of measure for the size of the CBS.

mbs

Syntax

mbs *size* [**kbits** | **bytes** | **kbytes**]

no mbs

Context

config>qos>network>ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command overrides the default MBS for the meter. The maximum burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the MBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter-configured parameters.



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

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

Default

512 kbits

Parameters

size

Specifies the number of kilobits or kilobytes or bytes reserved for the meter. For example, if a value of 100 kbits is desired, then enter the value 100. The bucket size is rounded off to the next highest 4096 bytes boundary.

Values **kbits** — 4 to 2146959, default
 bytes — 512 to 274810752
 kbytes — 1 to 268369

[kbits | bytes | kbytes]

Specifies the unit of measure for the size of the MBS.

mode

Syntax

mode *mode*

no mode

Context

config>qos>network>ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command defines the mode of the meter. The mode can be configured as Two Rate Three Color Marker (trTCM) or Single Rate Three Color Marker (srTCM). The **mode** command can be executed at any time.

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

Default

trtcm

Parameters

trtcm1

Keyword to meter the packet stream and mark the packets green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked yellow or green depending on whether it exceeds or does not exceed the CIR. The trTCM1 is useful for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate.

srtcm

Keyword to meter a packet stream and mark its packets green, yellow, or red. Marking is based on a CIR and two associated burst sizes, a CBS and an MBS. A packet is marked green if it does not exceed the CBS, yellow if it exceeds the CBS but not the CIR, and red otherwise. The srTCM is useful for ingress policing of a service, where only the length, not the peak rate, of the burst determines service eligibility.

rate

Syntax

rate *cir* *cir-rate-in-kbps* [**pir** *pir-rate-in-kbps*]

no rate

Context

config>qos>network>ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command defines the administrative PIR and CIR parameters for the meter.

The **rate** command can be executed at any time, altering the PIR and CIR for all meters created through the association of the Network QoS policy with the *meter-id*.

The **no** form of this command reverts all meter instances created with this *meter-id* to the default PIR (max) and CIR (0) parameters.



Note:

The value of rates are represented as 1000 kilobits per second and bursts are represented as 1024 kilobits per second.

Default

rate 0 pir max

Parameters

cir cir-rate-in-kbps

Specifies the administrative CIR, in kilobits, for the meter. The **cir** parameter overrides the default administrative CIR used by the meter. 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.

The actual CIR is dependent on the meter's **adaptation-rule** parameters and the hardware.

Values 0 to 20000000, max

pir pir-rate-in-kbps

Specifies the administrative PIR, in kilobits, for the meter. When this command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of max is assumed. When the **rate** command is executed, a PIR setting is optional.

The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the pir-rate value.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR is dependent on the meter's adaptation-rule parameters and the hardware.



Note:

If the meter mode is configured as **trtcm2**, the system configures the policer EIR rate, based on the value of the PIR rate configured by the user.

Values 0 to 20000000, max

7.6.2.5 Network egress QoS policy commands

egress

Syntax

egress

Context

config>qos>network policy-id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates or edits egress policy entries that specify the forwarding class map to be instantiated when this policy is applied to the network IP interface, or access-uplink port

The forwarding class and profile state mapping to appropriate marking values for all packets are defined in this context.

In network mode of operation, the system supports use of forwarding class mapping to EXP bits for IP interface, forwarding class mapping to DSCP and dot1p bits for network ports. In access-uplink mode of operation it allows the user to specify the FC mapping to dot1p bits for access-uplink ports.

fc

Syntax

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

Context

config>qos>network>egress

Platforms

7210 SAS-T (access-uplink mode)

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 marking criteria of packets flowing through it. The **fc** command overrides the default

parameters for that forwarding class to the values defined in the network default policy. Appropriate default parameters are picked up based on whether the **network-policy-type** is **port** or **ip-interface**.

The **no** form of this command removes the forwarding class LSP EXP/dot1p/DSCP map associated with this *fc*, as appropriate. The forwarding class reverts to the defined parameters in the default network policy. If the *fc-name* is removed from the network policy that forwarding class reverts to the factory defaults.

Default

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

Parameters

fc-name

Specifies a case-sensitive, system-defined forwarding class name for which policy entries will be created.

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

de-mark

Syntax

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

Context

config>qos>network>egress>fc

Platforms

7210 SAS-T (access-uplink mode)

Description

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

If no *de-value* is present, the default values are used for the marking of the DEI bit, as defined in the IEEE 802.1ad-2005 standard. For example 0 for in-profile packets, 1 for out-of-profile ones.

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

Parameters

de-value

Specifies the DEI value to use for this forwarding class.

Values 0 or 1

dot1p

Syntax

[no] **dot1p** *dot1p-value*

Context

config>qos>network>egress>fc

Platforms

7210 SAS-T access-uplink mode

Description

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the **dot1p** command has no effect.

DEI marking can be enabled using the **de-mark** command along with this command for the command to take effect. When **de-mark** command is configured along with this command, then the DEI bit is marked in the packet to indicate the profile of the packet. The DEI bit is marked to 0 to indicate in-profile/green packet and 1 to indicate out-of-profile/yellow packet. If the **force de-value** parameter is specified then the DEI bit is set to specified value for all packets.

If the **no** form of this command is executed then software will use the dot1p-in-profile and dot1p-out-profile if configured, else it will use default values.



Note:

The following rules are applied to determine the dot1p values when both the **dot1p** command, and the **dot1p-in-profile** and **dot1p-out-profile** commands are specified:

- If de-mark is not configured, **dot1p [in | out]-profile** values are considered. Even if **dot1p** value is configured it is ignored and if **dot1p [in | out]-profile** value is not configured, default values are considered for that FC.
- If de-mark is configured and the **dot1p** value is configured, it is considered. If the **dot1p [in | out]-profile** value is configured, it is considered. In this case, the **dot1p** value takes precedence over the **dot1p [in | out]-profile**.

Default

no dot1p

Parameters

dot1p-value

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

dot1p-in-profile

Syntax

```
dot1p-in-profile dot1p-priority  
no dot1p-in-profile
```

Context

```
config>qos>network>egress>fc
```

Platforms

7210 SAS-T access-uplink mode

Description

The command adds the capability to mark on an egress the in and out of profile status through a certain dot1p combination, similarly with the DEI options. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with in-profile status (or green color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined dot1p-value. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the dot1p command has no effect.

If DEI marking is enabled using the de-mark command and the command **dot1p dot1p-value** is used to configure the dot1p value, then this command has no effect. In other words, enabling DEI marking has precedence over this command and the system ignores this command.

When this command is used the DEI Bit is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DEI bit is set to 0.

The **no** form of this command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.



Note:

The following rules are applied to determine the dot1p values when both the **dot1p** command and the **dot1p-in-profile** and **dot1p-out-profile** commands are specified:

1. If **de-mark** is not configured, **dot1p [in|out]-profile** values are considered. Even if **dot1p** command value is configured, it is ignored, and if the **dot1p [in|out]-profile** value is not configured, default values are considered for that FC.
2. If de-mark is configured and if the **dot1p** command value is configured, it is considered. Otherwise, if the **dot1p [in|out]-profile** value is configured, it is considered. In this case, the **dot1p** command value has the precedence over the **dot1p [in|out]-profile**.
3. If marking is enabled and both the **dot1p** command value and the **dot1p-[in|out]-profile** commands are not specified, the default values are used.

Default

0

Parameters

dot1p-priority

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

dot1p-out-profile

Syntax

dot1p-out-profile *dot1p-priority*

no dot1p-out-profile

Context

config>qos>network>egress>fc

Platforms

7210 SAS-T access-uplink mode

Description

The command adds the capability to mark on an egress the in and out of profile status via a certain dot1p combination, similarly with the DEI options. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with out-of-profile status (or yellow color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined dot1p-value. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the dot1p command has no effect.

If DEI marking is enabled using the de-mark command and the dot1p-value is configured, then this command has no effect. In other words, enabling DEI marking has precedence over this command and the system ignores this command.

When this command is used the DEI Bit is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DEI bit is set to 0.

The **no** form of this command reverts the IEEE 802.1P or IEEE 802.1Q priority bits to the default.



Note:

The following rules are applied to determine the **dot1p** values to when both the **dot1p** command and **dot1p-in-profile** and **dot1p-out-profile** commands are specified:

1. If de-mark is not configured, the **dot1p [in|out]-profile** values are considered. Even if the **dot1p** command value is configured, it is ignored, and if the **dot1p [in|out]-profile** value is not configured, default values are considered for that FC.
2. If de-mark is configured and if the **dot1p** command value is configured, it is considered. Otherwise, if the **dot1p [in|out]-profile** value is configured, it is considered. In this case, the **dot1p** command value, has the precedence over the **dot1p [in|out]-profile** value.
3. If marking is enabled and the **dot1p** command value and the **dot1p-[in|out]-profile** command are not specified, the default values are used.

Default

0

Parameters

dot1p-priority

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

dscp-in-profile

Syntax

dscp-in-profile *dscp-name*

no dscp-in-profile

Context

config>qos>network >egress>fc

Platforms

7210 SAS-T (access-uplink mode)

Description

This command specifies the in-profile DSCP name for the forwarding class.

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

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

Parameters

dscp-name

Specifies a system- or user-defined, case-sensitive *dscp-name*.

dscp-out-profile

Syntax

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

Context

config>qos>network>egress>fc

Platforms

7210 SAS-T (access-uplink mode)

Description

This command specifies the out-of-profile DSCP name for the forwarding class.

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

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

Parameters

dscp-name

Specifies a system- or user-defined, case-sensitive *dscp-name*.

remarking

Syntax

remarking {**use-dot1p** | **use-dscp** | **all**}
no remarking

Context

config>qos>network>egress

Platforms

7210 SAS-T(access-uplink mode)

Description

This command enables the system to remark egress packets sent out of access-uplink ports. The user can specify if dot1p or dscp or both dot1p and dscp to be used for marking the packets sent out of the port.

When 7210 is operated in access-uplink mode, marking support is available as given below.

- On access-uplink port egress, the behavior is as follows:

- If the **use-dot1p** keyword is configured in the access-egress policy, then the dot1p bits are marked in the packet header for all traffic sent out of all SAPs configured on that access-uplink port.
- If the **use-dscp** keyword is configured in the access-egress policy, then the IP DSCP bits are marked in the packet header for IPv4 traffic sent out of all SAPs configured on that access port.



Note:

DSCP marking also marks the IPv4 packets associated with SAPs configured in an Layer 2 VPN service. If this is not required, to avoid this it is recommended to use only dot1p marking on access-uplink ports.

- If the **all** keyword is configured in the access-egress policy, then the dot1p bits are marked in the packet header for all traffic (Layer 2 and IPv4) sent out of all SAPs and the IP DSCP bits are marked in the packet header for all IPv4 traffic sent out of all SAPs configured on that access port.



Note:

- DSCP marking also marks the packets associated with SAPs configured in an Layer 2 VPN service. If this is not required, to avoid this it is recommended to use only dot1p marking on access-uplink ports.
- If remarking is enabled in access-egress policy, by default **use-dot1p** is used. If no marking values are specified, the default FC-to-dot1p marking values are used.

The **no** form of this command disables remarking.

Default

no remarking

remark

Syntax

remark *policy-id*

no remark

Context

config>qos>network>egress

Platforms

7210 SAS-T (network mode), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, 7210 SAS-Sx 10/100GE, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command specifies the remarking policy ID to use for marking packets on network port and network IP interface egress.

The remarking policy ID must be associated with the appropriate network QoS policy associated with the network port and or network IP interface and remarking must be enabled in the network QoS policy

to enable marking of packets sent out of network port and or network IP interface egress. Remarking policy of type **dot1p**, **dscp**, and **dot1p-dscp** is allowed to be used when the remark policy is associated with network QoS policy of type **port**. Remarking policy of type **lsp-exp** and **dot1p-lsp-exp-shared** is allowed to be used when the remark policy is associated with network QoS policy of type **ip-interface**. [Table 78: Summary of remark policy and attachment points for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE \(network mode\)](#) and [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) describe the remark policies, their attachment points supported on the node, and their uses.

The **no** form of this command removes the explicit association of the remark policy and associates the default remark policy. For example, if remarking is enabled and the **no remark** command is configured, the default remark policy is used to mark packets sent out. If no remark policy is executed and remarking is disabled, then packets are not remarked at all.

Parameters

policy-id

Specifies the remark policy ID.

Values 1 to 65535

remarking

Syntax

[no] remarking

Context

config>qos>network>egress

Platforms

7210 SAS-T (network mode), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, 7210 SAS-Sx 10/100GE, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command enables the system to remark egress packets sent out of network ports and hybrid ports.

When remarking is enabled, the remark policy configured in the QoS policy context is used to determine the FC to QoS bit mapping. For example, when remarking is enabled in the network QoS policy, the remark policy associated with network QoS policy is used to determine the FC-to-EXP mapping to use for marking packets sent out of access ports.

See [Remark policies](#) for the remark policy that can be used to configure FC to priority bit markings in different QoS policies associated with different service entities. For more information, see [Table 78: Summary of remark policy and attachment points for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE \(network mode\)](#) and [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#).

The **no** form of this command disables remarking. When remarking is disabled for MPLS LSR traffic, EXP values received at ingress are not modified at egress. For MPLS LER traffic, where the node adds the

MPLS encapsulation, MPLS EXP bits are set based on the mapping specified in the policy associated with the IP interface. If the user does not attach an explicit policy, the default policy is used.

Default

no remarking

scope

Syntax

scope {exclusive | template}

no scope

Context

config>qos>network policy-id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the network policy scope as exclusive or template.

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

Default

template

Parameters

exclusive

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

template

Specifies that the policy can be applied to multiple interfaces on the router.

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

7.6.2.6 Self-generated traffic commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)

sgt-qos

Syntax

sgt-qos

Context

config>router

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

Commands in this context configure DSCP/dot1p re-marking for select self-generated traffic.

application

Syntax

application *dscp-app-name* **dscp** {*dscp-value* | *dscp-name*}

application *dot1p-app-name* **dot1p** *dot1p-priority*

no application {*dscp-app-name* | *dot1p-app-name*}

Context

config>router>sgt-qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures DSCP/dot1p re-marking for self-generated application traffic. When an application is configured using this command, then the specified DSCP name/value is used for all packets generated by this application within the router instance it is configured. The instances can be base router or VPRN service.

Using the value configured in this command does the following:

- Sets the DSCP bits in the IP packet.
- Maps to the FC.
- Based on this FC, the egress QoS policy sets the Ethernet 802.1p and MPLS EXP bits. This includes ARP and IS-IS packets that, because of their nature, do not carry DSCP bits.
- The DSCP value in the egress IP header will be as configured in this command.

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

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

Parameters

dscp-app-name

Specifies the DSCP application name.

Values The following values apply to the base router instance:
bgp, dns, ftp, icmp, igmp, ldp, ndis, ntp, ospf, pim, ptp, radius, rip, rsvp,
snmp, snmp-notification, ssh, syslog, tacplus, telnet, tftp, traceroute,
vrrp, arp, isis

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 none, or 0 to 7

dot1p-app-name

Specifies the dot1p application name.

Values arp, isis

dscp

Syntax

dscp *dscp-name* **fc** *fc-name*

no dscp *dscp-name*

Context

config>router>sgt-qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

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

Self-generated traffic for configured applications 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 a 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. A 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. Applications and protocols that are configured under the **dscp** command will use the configured IP DSCP value.

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

7.6.2.7 Operational commands

copy

Syntax

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

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters

network *src-pol dst-pol*

Specifies the source policy that the copy command will copy and the destination policy to which the command will duplicate the policy to a new or different policy ID. This parameter indicates that the source and destination policies are network policy IDs.

Values 1 to 65535

overwrite

Keyword 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.

Example:

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

7.6.2.8 Show commands

- [Show commands \(7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12\)](#)

network

Syntax

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

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

Displays information about ingress and egress EXP bit mappings and network policy interface associations (for devices operating in network mode).

detail

Displays information about ingress and egress dot1p bit mappings and network policy interface associations (for 7210 SAS-T in access-uplink mode).

Output

The following outputs are examples of QoS network policy information, and the associated tables describe the output fields:

- [Sample output, Table 53: Output fields: QoS network policy](#)
- [Sample output \(policy ID 1\) for the 7210 SAS-R6 and 7210 SAS-R12, Table 54: Output fields: network QoS policyfor the 7210 SAS-R6 and 7210 SAS-R12](#)
- [Sample output \(policy ID 2\) for the 7210 SAS-R6 and 7210 SAS-R12, Table 54: Output fields: network QoS policyfor the 7210 SAS-R6 and 7210 SAS-R12](#)

Sample output

```
A:qos1# show qos network
=====
Network Policies
=====
Policy-Id      Remark  LerUseDscp  Description
-----
1              False   False      Default network-port QoS policy.
2              False   False      Default network QoS policy.
=====
A:qos1#

*A:ALA# show qos network 1 detail
=====
QoS Network Policy
=====
Network Policy (1)
-----
Policy-id      : 1              Remark      : False
Forward Class  : be              Profile     : Out
Attach Mode    : l2              Config Mode  : l2+mpls
Scope          : Template      Policy Type  : port
Accounting     : packet-based
Description    : Default network-port QoS policy.
-----
DSCP           Forwarding Class  Profile
-----
be             be              Out
ef             ef              In
cs1            l2              In
nc1            h1              In
nc2            nc              In
af11           af              In
af12           af              Out
```

```

af41                                     h2                                     In
-----
Dot1p Bit Map                          Forwarding Class                      Profile
-----
No Matching Entries
-----
Meter Mode   CIR Admin   CIR Rule   PIR Admin   PIR Rule   CBS       MBS
-----
1    TrTcm_CA  0           closest    max        closest   32 KBytes 128 KBytes
-----
FC           UCastM      MCastM
-----
No FC-Map Entries Found.
-----
Egress Forwarding Class Queuing
-----
FC Value      : 0                      FC Name       : be
- DSCP Mapping
Out-of-Profile : be                      In-Profile    : be

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

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

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

FC Value      : 2                      FC Name       : af
- DSCP Mapping
Out-of-Profile : af12                   In-Profile    : af11

- Dot1p Mapping
Out-of-Profile : 2                      In-Profile    : 3

FC Value      : 3                      FC Name       : l1
- DSCP Mapping
Out-of-Profile : af22                   In-Profile    : af21

- Dot1p Mapping
Out-of-Profile : 2                      In-Profile    : 3

FC Value      : 4                      FC Name       : h2
- DSCP Mapping
Out-of-Profile : af41                   In-Profile    : af41

- Dot1p Mapping
Out-of-Profile : 4                      In-Profile    : 4

FC Value      : 5                      FC Name       : ef
- DSCP Mapping
Out-of-Profile : ef                     In-Profile    : ef

- Dot1p Mapping
Out-of-Profile : 5                      In-Profile    : 5

FC Value      : 6                      FC Name       : h1
- DSCP Mapping
Out-of-Profile : nc1                    In-Profile    : nc1

- Dot1p Mapping
Out-of-Profile : 6                      In-Profile    : 6

```

```
FC Value      : 7
- DSCP Mapping
Out-of-Profile : nc2
FC Name       : nc
In-Profile    : nc2

- Dot1p Mapping
Out-of-Profile : 7
In-Profile     : 7
-----
Interface Association
-----
No Interface Association Found.
-----
Port Attachments
-----
Port-id : 1/1/1
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
Port-id : 1/1/9
Port-id : 1/1/10
Port-id : 1/1/11
Port-id : 1/1/12
Port-id : 1/1/13
Port-id : 1/1/14
Port-id : 1/1/16
Port-id : 1/1/17
Port-id : 1/1/18
Port-id : 1/1/20
Port-id : 1/1/21
Port-id : 1/1/22
Port-id : 1/1/23
Port-id : 1/1/24
=====
*A:ALA#

*A:ALA# show qos network 2 detail
=====
QoS Network Policy
=====
Network Policy (2)
-----
Policy-id      : 2
Forward Class  : be
Attach Mode    : mpls
Scope          : Template
Accounting     : packet-based
Profile Policy : 1
Global Prof    : 1
Description    : Default network QoS policy.

-----
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
7
h1
nc
In
In
-----
Meter Mode    CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS      MBS
-----
1    TrTcm_CA  0         closest    max        closest   32 KBytes 128 KBytes
9    TrTcm_CA  0         closest    max        closest   32 KBytes 128 KBytes
-----
FC              UCastM      MCastM
-----
No FC-Map Entries Found.
-----
Egress Forwarding Class Queuing
-----
FC Value       : 0                FC Name       : be
- LSP EXP Bit Mapping
Out-of-Profile : 0                In-Profile    : 0

FC Value       : 1                FC Name       : l2
- LSP EXP Bit Mapping
Out-of-Profile : 1                In-Profile    : 1

FC Value       : 2                FC Name       : af
- LSP EXP Bit Mapping
Out-of-Profile : 2                In-Profile    : 3

FC Value       : 3                FC Name       : l1
- LSP EXP Bit Mapping
Out-of-Profile : 2                In-Profile    : 3

FC Value       : 4                FC Name       : h2
- LSP EXP Bit Mapping
Out-of-Profile : 4                In-Profile    : 4

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

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

FC Value       : 7                FC Name       : nc
- LSP EXP Bit Mapping
Out-of-Profile : 7                In-Profile    : 7
-----
Interface Association
-----
Interface      : system
IP Addr.       : n/a                Port Id       : system
Interface      : in-band-management
IP Addr.       : 10.135.25.189/24    Port Id       : 1/1/23
-----
Port Attachments
-----
No Matching Entries
=====
*A:ALA#

For SAS-MXP:
*A:qos1# show qos network 1001 detail
=====
QoS Network Policy

```

Network Policy (1001)													
Policy-id : 1001						Remark : False							
Forward Class : be						Profile : In							
Attach Mode : mpls						Config Mode : mpls							
Scope : Template						Policy Type : IpInterface							
Accounting : packet-based													
Description : ip-interface-type													
LSP EXP Bit Map						Forwarding Class				Profile			
0						be				Out			
1						l2				Out			
2						af				In			
3						l1				Out			
4						h2				In			
5						ef				Out			
6						h1				Out			
7						nc				In			
Meter	Mode	CIR	Admin	CIR	Rule	PIR	Admin	PIR	Rule	CBS	Admin	MBS	Admin
			CIR			PIR		PIR		CBS	Oper	MBS	Oper
1	TrTcm_CA	4000		closest		8000		closest		def		def	
		4000				8000				def		500	
2	TrTcm_CA	4000		closest		7000		closest		16384		16384	
		4000				7000				16000		16000	
3	TrTcm_CA	4000		closest		7000		closest		def		def	
		4000				7000				def		500	
4	TrTcm_CA	4000		closest		7000		closest		def		def	
		4000				7000				def		500	
5	TrTcm_CA	4000		closest		7000		closest		def		def	
		4000				7000				def		500	
6	TrTcm_CA	4000		closest		7000		closest		def		def	
		4000				7000				def		500	
7	TrTcm_CA	4000		closest		7000		closest		def		def	
		4000				7000				def		500	
8	TrTcm_CA	7000		closest		7000		closest		def		def	
		7000				7000				def		500	
9	TrTcm_CA	4000		closest		7000		closest		def		def	
		4000				7000				def		500	
10	TrTcm_CA	4000		closest		7000		closest		def		def	
		4000				7000				def		500	
11	TrTcm_CA	4000		closest		7000		closest		def		def	
		4000				7000				def		500	
12	TrTcm_CA	4000		closest		7000		closest		def		def	
		4000				7000				def		500	
FC		UCastM				MCastM							
l2		2				def							
af		3				def							
l1		4				def							
h2		5				12							
ef		6				11							
h1		7				10							
nc		8				9							
Egress Forwarding Class Queuing													
FC Value : 0						FC Name : be							
- LSP EXP Bit Mapping													

```

Out-of-Profile : 0                      In-Profile   : 0
FC Value       : 1                      FC Name      : l2
- LSP EXP Bit Mapping
Out-of-Profile : 1
...
=====
*A:qos1#

```

Table 53: Output fields: QoS network policy

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 EXP bit mapping defined under the egress node of the network QoS policy.
Description	A text string that helps identify the policy's context in the configuration file
Forward Class/ FC Name	Specifies the forwarding class name
Profile	Out— Specifies the EXP marking for the packets which are out-of-profile, egressing on this queue Specifies the EXP marking for the packets which are out-of-profile, egressing on this queue. In — Specifies the EXP marking for the packets which are in-of-profile, egressing on this queue Specifies the EXP markings for in-profile packets egressing this queue
Accounting	Packet-based — Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow Frame-based — Specifies that the meters associated with this policy account for the packet framing overheads (such as for Ethernet the IFG and preamble), while accounting the bandwidth to be used by the flow
Profile policy	Displays the profile policy ID
Global Prof	Displays the global profile policy ID for LDP packets
EXP Bit Mapping:	
Out-of-Profile	Displays the EXP value used for out-of-profile traffic

Label	Description
In-Profile	Displays the 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

Sample output (policy ID 1) for the 7210 SAS-R6 and 7210 SAS-R12

```
*A:SAS>config>qos>network# show qos network 1 detail

=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id       : 1
Egr Remark      : False                      Egr Rem Plcy : N/A
Forward Class   : be                        Profile      : Out
Scope           : Template                  Policy Type   : port
Accounting      : packet-based
Description     : Default network-port QoS policy.

-----
Dot1p Bit Map           Forwarding Class           Profile
-----
No Matching Entries

-----
Meter Mode      CIR Admin CIR Rule   PIR Admin   PIR Rule   CBS Admin MBS Admin
                CIR Oper                PIR Oper                CBS Oper  MBS Oper
-----
1      TrTcm1_CA    0         closest      max         closest    def        def
                0                                max
-----

FC           UCastM           MCastM
-----
No FC-Map Entries Found.

-----
Port Attachments
-----
Port-id : 1/1/10
Port-id : 1/1/11
Port-id : 1/1/12
Port-id : 1/1/13
Port-id : 1/1/20
Port-id : 1/1/21
Port-id : 1/1/22
Port-id : 1/1/23
Port-id : 1/1/24
Port-id : 1/1/25
Port-id : 1/1/26
Port-id : lag-3
Port-id : lag-5
```

```
=====
*A:SAS>config>qos>network#
```

Sample output (policy ID 2) for the 7210 SAS-R6 and 7210 SAS-R12

```
*A:SAS>config>qos>network# show qos network 2 detail
```

```
=====
QoS Network Policy
=====
```

```
-----
Network Policy (2)
-----
```

```
Policy-id      : 2
Egr Remark     : False                      Egr Rem Plcy : N/A
Forward Class  : be                        Profile      : Out
Scope          : Template                  Policy Type   : IpInterface
Accounting     : packet-based
Profile Policy  : 1
Local FC       : Disabled                  Global Prof   : 1
Description    : Default network QoS policy.
```

```
-----
Dot1p Bit Map          Forwarding Class          Profile
-----
```

```
No Matching Entries
```

```
-----
Meter Mode      CIR Admin CIR Rule   PIR Admin   PIR Rule    CBS Admin MBS Admin
                CIR Oper                PIR Oper
-----
1      TrTcm1_CA  0          closest    max         closest    7 KBytes 10 KBytes
                0
9      TrTcm1_CA  0          closest    max         closest    11 KBytes 14 KBytes
```

```
-----
FC          UCastM          MCastM
-----
```

```
No FC-Map Entries Found.
```

```
-----
Interface Association
-----
```

```
Interface      : system
IP Addr.       : 180.10.10.10/32          Port Id        : system
Interface      : in-band-management
IP Addr.       : 10.135.25.183/24         Port Id        : 1/1/24
Interface      : management
IP Addr.       : 10.135.25.183/24         Port Id        : A/1
```

```
=====
*A:SAS>config>qos>network#
```

```
*A:qos1# show qos network 1001 detail
```

```
=====
QoS Network Policy
=====
```

```
-----
Network Policy (1001)
-----
```

Policy-id	: 1001	Remark	: False
Forward Class	: be	Profile	: In
Attach Mode	: mpls	Config Mode	: mpls
Scope	: Template	Policy Type	: IpInterface
Accounting	: packet-based		
Description	: ip-interface-type		

LSP EXP Bit Map		Forwarding Class	Profile

0		be	Out
1		l2	Out
2		af	In
3		l1	Out
4		h2	In
5		ef	Out
6		h1	Out
7		nc	In

Meter Mode	CIR Admin	CIR Rule	PIR Admin
	CIR Oper		PIR Oper
		PIR Rule	CBS Admin
			CBS Oper
			MBS Admin
			MBS Oper

1	TrTcm_CA	4000	closest
		8000	closest
		def	def
2	TrTcm_CA	4000	closest
		7000	closest
		16384	16384
		4000	16000
3	TrTcm_CA	4000	closest
		7000	closest
		def	def
		4000	500
4	TrTcm_CA	4000	closest
		7000	closest
		def	def
		4000	500
5	TrTcm_CA	4000	closest
		7000	closest
		def	def
		4000	500
6	TrTcm_CA	4000	closest
		7000	closest
		def	def
		4000	500
7	TrTcm_CA	4000	closest
		7000	closest
		def	def
		4000	500
8	TrTcm_CA	7000	closest
		7000	closest
		def	def
		7000	500
9	TrTcm_CA	4000	closest
		7000	closest
		def	def
		4000	500
10	TrTcm_CA	4000	closest
		7000	closest
		def	def
		4000	500
11	TrTcm_CA	4000	closest
		7000	closest
		def	def
		4000	500
12	TrTcm_CA	4000	closest
		7000	closest
		def	def
		4000	500

FC	UCastM	MCastM	

l2	2	def	
af	3	def	
l1	4	def	
h2	5	12	
ef	6	11	
h1	7	10	
nc	8	9	

Egress Forwarding Class Queuing			

FC Value	: 0	FC Name	: be
- LSP EXP Bit Mapping		In-Profile	: 0
Out-of-Profile	: 0		
FC Value	: 1	FC Name	: l2

```
- LSP EXP Bit Mapping
Out-of-Profile : 1
...
=====
*A:qos1#
```

Table 54: Output fields: network QoS policy for the 7210 SAS-R6 and 7210 SAS-R12

Label	Description
Policy-Id	Displays 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 FC to bit mapping defined under the egress node of the network QoS policy.
Description	Displays a text string that helps identify the policy context in the configuration file
Forward Class/FC Name	Specifies the FC name
Profile	Out — Specifies the EXP marking for the packets that are out-of-profile, egressing on this queue In — Specifies the EXP marking for the packets that are in-of-profile, egressing on this queue
Accounting	Packet-based — Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow Frame-based — Specifies that the meters associated with this policy account for the packet framing overheads (such as for Ethernet the IFG and preamble), while accounting for the bandwidth used by the flow
Profile policy	Displays the profile policy ID
Global Prof	Displays the global profile policy ID for LDP packets
Bit Mapping:	
Out-of-Profile	Displays the value used for out-of-profile traffic
In-Profile	Displays the 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

mpls-lsp-exp-profile-map

Syntax

mpls-lsp-exp-profile-map [*policy-id*] [**detail**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays profile policy information.

Parameters

policy-id

Displays information for the specific policy ID.

Values 1 to 65535

detail

Displays detailed policy information.

Output

The following output is an example of MPLS LSP EXP profile mapping information, and [Table 55: Output fields: QoS MPLS LSP EXP profile map](#) describes the output fields.

Sample output

```
*A:7210-SAS>show>qos# mpls-lsp-exp-profile-map 1

=====
QoS MPLS LSP EXP Profile Maps
=====
-----
Profile Map-id      : 1
Description         : Default MPLS LSP EXP Profile Map policy
-----

Exp      Profile
-----
0        Out
1        In
2        Out
3        In
4        In
5        In
6        In
7        In
=====

*A:7210SAS>show>qos# mpls-lsp-exp-profile-map 1 detail
```

```
=====
QoS MPLS LSP EXP Profile Maps
=====
-----
Profile Map-id      : 1
Description         : Default MPLS LSP EXP Profile Map policy
-----
Exp      Profile
-----
0        Out
1        In
2        Out
3        In
4        In
5        In
6        In
7        In
-----
Network Policy Associations
-----
Network Policy Id   : 2
-----
=====
*A: 7210-SAS>show>qos#
```

Table 55: Output fields: QoS MPLS LSP EXP profile map

Label	Description
Profile Map-id	Displays the profile Map ID
Description	A text string that helps identify the policy's context in the configuration file
Exp	Displays the EXP. value
Profile	Specifies the marking of the packets as in-profile or out-of-profile
Network Policy Id	Displays the Network policy ID with which the mpls-lsp-exp-profile is associated

7.6.2.9 Show commands (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)

dscp-table

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

Context
show>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays DSCP name and DSCP value mappings.

Parameters

- value dscp-value

Specifies the DSCP value for which to display information.
- Values

0 to 63
- Default

all values

Output

The following output is an example of DSCP value information, and [Table 56: Output fields: QoS network DSCP](#) 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 56: Output fields: QoS network DSCP

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

router

Syntax

```
router [router-instance]  
router service-name service-name
```

Context

show

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays router information.

Parameters

<i>router-instance</i>	
Specifies the router name or service ID	
Values	<i>router-name</i> — Base <i>service-id</i> — 1 to 2147483647
Default	Base
<i>service-name</i>	
Specifies the service name, up to a 64 characters.	

sgt-qos

Syntax

```
sgt-qos
```

Context

show>router

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays QoS information about self-generated traffic. In the output, the value "none" indicates that the default value is used; it does not indicated that there is no value set. [Table 50: Default DSCP mapping table](#) lists the application defaults.

application

Syntax

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

Context

show>router>sgt-qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays application QoS settings.

Parameters

app-name

Specifies the application.

Values The following values apply to the base router instance:
bgp, dns, ftp, icmp, igmp, ldp, ndis, ntp, ospf, pim, ptp, radius, rip, rsvp,
snmp, snmp-notification, ssh, syslog, tacplus, telnet, tftp, traceroute,
vrrp, arp, isis
The following values apply to a VPRN service instance:
bgp, icmp, igmp, ndis, ospf, pim, ssh, telnet, traceroute, vrrp, arp



Note:

- The value "ptp" in the context of SGT QoS is defined as Precision Timing Protocol and is an application. The PTP application name is also used in areas such as event-control and logging. Precision Timing Protocol is defined in IEEE 1588-2008.
- The value "ptp" in the context of IP filters is defined as Performance Transparency Protocol. IP protocols can be used as IP filter match criteria; the match is made on the 8-bit protocol field in the IP header.

dscp

Displays all DSCP applications.

dot1p

Displays all dot1p applications.

Output

The following outputs show an example of application QoS information for the base router and for a VPRN service instance, and [Table 57: Output fields: SGT-QoS application](#) describes the output fields.

- [Sample output \(router\)](#)
- [Sample output \(VPRN service instance\)](#)

Sample output (router)

```
*A:SAS-DUTA# show router sgt-qos application
=====
DSCP Application Values
=====
Application          DSCP Value          Default DSCP Value
-----
bgp                   none                 none
dns                   none                 none
ftp                   none                 none
icmp                  none                 none
igmp                  none                 none
ldp                   none                 none
ndis                  none                 none
ntp                   none                 none
ospf                  none                 none
pim                   none                 none
ptp                   none                 none
radius                none                 none
rip                   none                 none
rsvp                  none                 none
snmp                  none                 none
snmp-notification    none                 none
ssh                   none                 none
syslog                none                 none
tacplus               none                 none
telnet                none                 none
tftp                  none                 none
traceroute            none                 none
vrrp                  none                 none
=====

Dot1p Application Values
=====
Application          Dot1p Value          Default Dot1p Value
-----
arp                   none                 none
isis                  none                 none
=====

*A:SAS-DUTA#

*A:SAS-DUTA# show router sgt-qos application arp
=====
Dot1p Application Values
=====
Application          Dot1p Value          Default Dot1p Value
-----
arp                   none                 none
=====

*A:SAS-DUTA#
```

Sample output (VPRN service instance)

```
=====
*A:SAS-DUTA# show router 1 sgt-qos application
=====
DSCP Application Values
=====
Application          DSCP Value          Default DSCP Value
-----
bgp                   none                 none
icmp                  cp17                 none
igmp                  none                 none
ndis                  none                 none
ospf                  none                 none
pim                   none                 none
ssh                   none                 none
telnet                none                 none
traceroute            none                 none
vrrp                  none                 none
=====
Dot1p Application Values
=====
Application          Dot1p Value          Default Dot1p Value
-----
arp                   none                 none
isis                  none                 none
=====
*A:SAS-DUTA#
```

```
*A:SAS-DUTA>config>service# \show router 1 sgt-qos application arp
=====
Dot1p Application Values
=====
Application          Dot1p Value          Default Dot1p Value
-----
arp                   none                 none
=====
*A:SAS-DUTA#
```

Table 57: Output fields: SGT-QoS application

Label	Description
Application	The DSCP or dot1p application
DSCP Value	The DSCP name or value assigned to the application; if you assign a value to the application (0 to 63), the DSCP name that maps to the value is displayed
Default DSCP Value	The default DSCP value
Dot1p Value	The dot1p priority assigned to the application (applies only to ARP and IS-IS)
Default Dot1p Value	The default dot1p value

dscp-map

Syntax

dscp-map [*dscp-name*]

Context

show>router>sgt-qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays DSCP-to-FC mappings.

Parameters

dscp-name

Specifies the DSCP name.

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

Output

The following outputs show examples of DSCP-to-FC mapping information for a base router and a VPRN service instance, and [Table 58: Output fields: SGT-QoS DSCP-to-FC mapping](#) describes the output fields.

- [Sample output \(router\)](#)
- [Sample output \(VPRN service instance\)](#)

Sample output (router)

```
*A:SAS-DUTA# show router sgt-qos dscp-map
=====
DSCP to FC Mappings
=====
DSCP Value      FC Value      Default FC Value
-----
be              nc            nc
cp1             be            be
cp2             be            be
cp3             be            be
cp4             be            be
cp5             be            be
cp6             be            be
cp7             be            be
cs1             be            be
cp9             be            be
```

af11	af	af
cp11	be	be
af12	af	af
cp13	be	be
af13	af	af
cp15	be	be
cs2	be	be
cp17	be	be
af21	l1	l1
cp19	be	be
af22	l1	l1
cp21	be	be
af23	l1	l1
cp23	be	be
cs3	be	be
cp25	be	be
af31	l1	l1
cp27	be	be
af32	l1	l1
cp29	be	be
af33	h2	l1
cp31	be	be
cs4	be	be
cp33	be	be
af41	nc	nc
cp35	be	be
af42	h2	h2
cp37	be	be
af43	h2	h2
cp39	be	be
cs5	be	be
cp41	be	be
cp42	be	be
cp43	be	be
cp44	be	be
cp45	be	be
ef	ef	ef
cp47	be	be
nc1	nc	nc
cp49	be	be
cp50	h2	h2
cp51	be	be
cp52	be	be
cp53	be	be
cp54	be	be
cp55	be	be
nc2	nc	nc
cp57	be	be
cp58	be	be
cp59	be	be
cp60	be	be
cp61	be	be
cp62	be	be
cp63	be	be
=====		
*A: SAS-DUTA#		

Sample output (VPRN service instance)

```
*A:SAS-DUTA# show router 1 sgt-qos dscp-map
=====
DSCP to FC Mappings
=====
```

DSCP Value	FC Value	Default FC Value
be	nc	nc
cp1	be	be
cp2	be	be
cp3	be	be
cp4	be	be
cp5	be	be
cp6	be	be
cp7	be	be
cs1	be	be
cp9	be	be
af11	af	af
cp11	be	be
af12	af	af
cp13	be	be
af13	af	af
cp15	be	be
cs2	be	be
cp17	ef	be
af21	l1	l1
cp19	be	be
af22	l1	l1
cp21	be	be
af23	l1	l1
cp23	be	be
cs3	be	be
cp25	be	be
af31	l1	l1
cp27	be	be
af32	l1	l1
cp29	be	be
af33	l1	l1
cp31	be	be
cs4	be	be
cp33	be	be
af41	nc	nc
cp35	be	be
af42	h2	h2
cp37	be	be
af43	h2	h2
cp39	be	be
cs5	be	be
cp41	be	be
cp42	be	be
cp43	be	be
cp44	be	be
cp45	be	be
ef	ef	ef
cp47	be	be
nc1	nc	nc
cp49	be	be
cp50	h2	h2
cp51	be	be
cp52	be	be
cp53	be	be
cp54	be	be
cp55	be	be
nc2	nc	nc
cp57	be	be
cp58	be	be
cp59	be	be
cp60	be	be
cp61	be	be

cp62	be	be
cp63	be	be
=====		
*A: SAS-DUTA#		

Table 58: Output fields: SGT-QoS DSCP-to-FC mapping

Label	Description
DSCP Value	Displays the DSCP values (displayed as names) of the self-generated traffic
FC Value	Displays the FC value mapped to each DSCP value
Default FC Value	Displays the default FC value

8 Network queue QoS policies

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

8.1 Overview

Network queue policies define the network egress queue parameters. Network queue policies are associated with the following:

- network ports on 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp in network mode
- hybrid ports on 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp in network mode
- access uplink ports on 7210 SAS-T in access uplink mode

Network queue policies define the bandwidth distribution for the network egress queues.

There is one default network queue policy. 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 all network ports and hybrid port. The user must explicitly create and then associate other network queue QoS policies if the default values need to be modified.

The default network queue policy is defined as follows:

- On 7210 SAS-T (both access-uplink and network mode), it defines 8 forwarding classes and assigns 8 queues, one to each FC.
- On 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, it defines 8 forwarding classes and assigns 8 queues, one to each FC.
- On 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, it defines 8 forwarding classes and assigns 16 queues, one each to unicast traffic mapped to the FC and to multicast traffic mapped to the FC, for a total of 2 egress queues per FC.

8.2 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.

8.2.1 Create a network queue QoS policy

Configuring and applying QoS policies other than the default policy is optional. A default network queue policy is applied to all network ports (for 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T, 7210 SAS-Sx/S

1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp in network mode) and access uplink ports (for 7210 SAS-T in access uplink mode).

To create a network queue policy, define the following:

- **network queue policy name**

The system will not dynamically assign a name.

- **description**

The description provides a brief overview of policy features.

FCs are mapped to the queues according to [Table 31: Forwarding class to queue-ID map](#) and [Table 32: Forwarding class-to-queue ID map for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#)

Use the following syntax to create a network queue QoS policy.

```
config>qos
  network-queue policy-name
    description description-string
    queue queue-id
      rate cir cir-percent [pir pir-percent]
      adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
```

Output example

The following is a sample configuration output to create a network queue QoS policy on the 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T.

```
*A:Dut-B>config>qos>network-queue# info detail
-----
description "Default network queue QoS policy."
queue 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
exit
queue 2
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
exit
queue 3
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
exit
queue 4
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
exit
queue 5
  rate cir 100 pir 100
  adaptation-rule cir closest pir closest
exit
queue 6
  rate cir 100 pir 100
  adaptation-rule cir closest pir closest
exit
queue 7
  rate cir 10 pir 100
  adaptation-rule cir closest pir closest
exit
queue 8
  rate cir 10 pir 100
```

```
        adaptation-rule cir closest pir closest
        exit
-----
*A:Dut-B>config>qos>network-queue#
```

Output example

The following is a sample configuration output to create a network queue QoS policy on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

```
*A:7210SAS>config>qos# info detail
-----
#-----
echo "QoS Policy Configuration"
#-----
    network-queue "t" create
    no description
    queue 1
        rate cir 0 pir 100
        adaptation-rule cir closest pir closest
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
    queue 2
        rate cir 0 pir 100
        adaptation-rule cir closest pir closest
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
    queue 3
        rate cir 0 pir 100
        adaptation-rule cir closest pir closest
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
    queue 4
        rate cir 0 pir 100
        adaptation-rule cir closest pir closest
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
    queue 5
        rate cir 0 pir 100
        adaptation-rule cir closest pir closest
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
    queue 6
        rate cir 0 pir 100
        adaptation-rule cir closest pir closest
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
    queue 7
        rate cir 0 pir 100
```

```
        adaptation-rule cir closest pir closest
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
queue 8
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    queue-mgmt "default"
    queue-mode weighted
    weight 1
exit
-----
*A:7210SAS>config>qos
```

8.2.2 Applying network queue policies

Apply network queue policies to the following entities:

- Ethernet Ports

8.2.2.1 Applying network queue configuration in network mode

Use the following syntax to apply a network queue policy to an Ethernet port in network mode of operation.

The network-queue policy can only be applied on a network port.

```
config>port#
    ethernet
```

Example

```
#-----
echo "Port Configuration"
#-----
    port 1/1/1
        ethernet
            mode network
            network

                queue-policy "nql-cbs"
            exit
        exit
    exit
    no shutdown
exit
```

8.2.2.2 Applying network queue configuration in access-uplink mode

Use the following syntax to apply a network queue policy to an Ethernet port in access-uplink mode of operation.

```
config>port#
    ethernet
```

```
access
  uplink
  queue-policy policy-name
```

Example

```
#-----
echo "Port Configuration"
#-----
port 1/1/1
  ethernet
    mode access uplink
    access
      uplink
        queue-policy "nq1-cbs"
    exit
  exit
exit
no shutdown
exit
```

8.3 Default network queue policy values

The default network queue policies are identified as **policy-id default**. The default policies cannot be modified or deleted for 7210 SAS-T (in access-uplink and network modes), 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE devices.

Example

```
A:qos1# show qos network-queue default detail
=====
QoS Network Queue Policy
=====
Network Queue Policy (default)
-----
Policy          : default
Accounting      : packet-based
Description     : Default network queue QoS policy.

-----
Queue  CIR      PIR      CBS
      CIR Rule  PIR Rule
-----
1      0        100      12.50
      closest  closest
2      25        100      12.50
      closest  closest
3      25        100      12.50
      closest  closest
4      25        100      12.50
      closest  closest
5      100       100      12.50
      closest  closest
6      100       100      12.50
      closest  closest
7      10        100      12.50
      closest  closest
8      10        100      12.50
      closest  closest
```

```

-----
FC      UCastQ
-----
be      1
l2      2
af      3
l1      4
h2      5
ef      6
h1      7
nc      8

-----
Associations
-----
Port-id : 1/1/4
Port-id : 1/1/8
Port-id : 1/1/9
Port-id : 1/1/10
Port-id : 1/1/12
Port-id : 1/1/13
Port-id : 1/1/14
Port-id : 1/1/15
Port-id : 1/1/16
Port-id : 1/1/17
Port-id : 1/1/18
Port-id : 1/1/19
Port-id : 1/1/20
Port-id : 1/1/21
Port-id : 1/1/22
Port-id : 1/1/23
Port-id : 1/1/24
Port-id : lag-1
=====
A:qos1#

```

The following displays default policy parameters for 7210 SAS-T (access-uplink and network modes), 7210 SAS-S 1/10GE, and 7210 SAS-Sx 10/100GE.

Example

```

*A:Dut-C>config>qos>network-queue# info detail
-----
description "Default network queue QoS policy."
queue 1
    rate 0 pir 100
    adaptation-rule cir closest pir closest
exit
queue 2
    rate 25 pir 100
    adaptation-rule cir closest pir closest
exit
queue 3
    rate 25 pir 100
    adaptation-rule cir closest pir closest
exit
queue 4
    rate 25 pir 100
    adaptation-rule cir closest pir closest
exit
queue 5
    rate 100 pir 100

```

```
        adaptation-rule cir closest pir closest
    exit
    queue 6
        rate 100 pir 100
        adaptation-rule cir closest pir closest
    exit
    queue 7
        rate 10 pir 100
        adaptation-rule cir closest pir closest
    exit
    queue 8
        rate 10 pir 100
        adaptation-rule cir closest pir closest
    exit
-----
*A:Dut-C>config>qos>network-queue#
```

The following displays default policy parameters for the default network queue policy for 7210 SAS-Mxp.

Example

```
*A:sim_dutc>config>qos>network-queue$ info detail
-----
no description
queue 1
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    queue-mgmt "default"
    queue-mode weighted
    weight 1
exit
queue 2
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    queue-mgmt "default"
    queue-mode weighted
    weight 1
exit
queue 3
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    queue-mgmt "default"
    queue-mode weighted
    weight 1
exit
queue 4
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    queue-mgmt "default"
    queue-mode weighted
    weight 1
exit
queue 5
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    queue-mgmt "default"
    queue-mode weighted
    weight 1
exit
```

The following are sample default policy parameters for 7210 SAS-R6 and 7210 SAS-R12.

Example

```
*A:7210SAS>show>qos# network-queue "default" detail
```

QoS Network Queue Policy

Network Queue Policy (t)

```
Policy      : t
Accounting  : packet-based
Description : (Not Specified)
Wrr Policy  :
Pkt.Byte Offset: 0
```

Queue Rates and Rules

QueueId	CIR(%)	CIR Adpt Rule	PIR(%)	PIR Adpt Rule
Queue1	0	closest	100	closest
Queue2	0	closest	100	closest
Queue3	0	closest	100	closest
Queue4	0	closest	100	closest
Queue5	0	closest	100	closest
Queue6	0	closest	100	closest
Queue7	0	closest	100	closest
Queue8	0	closest	100	closest

Queue Mode and Weight Details

QueueId	Mode	Weight
Queue1	weighted	1
Queue2	weighted	1
Queue3	weighted	1
Queue4	weighted	1
Queue5	weighted	1
Queue6	weighted	1
Queue7	weighted	1
Queue8	weighted	1

High Slope

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	75
Queue2	Down	70	90	75
Queue3	Down	70	90	75
Queue4	Down	70	90	75
Queue5	Down	70	90	75
Queue6	Down	70	90	75
Queue7	Down	70	90	75
Queue8	Down	70	90	75

Low Slope

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	75


```

Queue2      Down      50      75      75
Queue3      Down      50      75      75
Queue4      Down      50      75      75
Queue5      Down      50      75      75
Queue6      Down      50      75      75
Queue7      Down      50      75      75
Queue8      Down      50      75      75
-----
Burst Sizes and Time Average Factor
-----
-----
QueueId      CBS      MBS      Time Average Factor      Queue-Mgmt
-----
Queue1      def      def      7      default
Queue2      def      def      7      default
Queue3      def      def      7      default
Queue4      def      def      7      default
Queue5      def      def      7      default
Queue6      def      def      7      default
Queue7      def      def      7      default
Queue8      def      def      7      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
-----
Network-Port Associations
-----
No Matching Entries
=====
*A:7210SAS>show>qos#

*7210SAS>config>qos>network-queue# info detail
-----
description "Default hybrid queue QoS policy."
queue 1
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 2
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 3
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit

```

```
queue 4
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 5
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 6
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 7
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 8
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
-----
*7210SAS>config>qos>network-queue#
```

8.4 Default network queue policy values for hybrid ports on 7210 SAS-R6 and 7210 SAS-R12

The default network queue policies are identified as *policy-id default* and *_tmnx_hybrid_default*.The MPLS and IP traffic on hybrid ports arrives at the *_tmnx_hybrid_default* policy if there is no policy applied.

Output example

The following is a sample configuration output for default policy parameters for 7210 SAS-R6 and 7210 SAS-R12.

```
*A: 7210SAS>show>qos# network-queue

=====
Network Queue Policies
=====
Policy-Id          Description
-----
default            Default network queue QoS policy.
_tmnx_hybrid_default  Default hybrid queue QoS policy.
=====
*A: 7210SAS>show>qos#
```

8.5 Service management tasks

This section describes the service management tasks.

8.5.1 Deleting network queue QoS policies

A network queue policy is associated by default with all network ports (for devices operating in network mode) and access uplink ports (for 7210 SAS-T in access-uplink mode). The user can replace the default policy with a customer-configured policy, but you cannot entirely remove a QoS policy. When you remove a QoS policy, the policy association reverts to the default network-queue policy **default**.

A network-queue policy cannot be deleted until it is removed from all network ports where it is applied.

Use the following syntax to delete a user-created network queue policy.

```
config>qos# no network-queue policy-name
```

Example:

```
config>qos# no network-queue nq1
```

8.5.2 Copying and overwriting network queue QoS policies

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

Use the following syntax to copy and overwrite a QoS policy.

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

Example:

```
config>qos# copy network-queue nq1-cbs nq2-cbs
```

Output example

The following is a sample of the copied policies output.

```
*A:card-1>config>qos# info
#-----
echo "QoS Slope and Queue Policies Configuration"
#-----
.....
    network-queue "nq1-cbs" create
        queue 1
            rate cir 0 pir 32
            adaptation-rule cir max
        exit
        queue 2
        exit
        queue 3
```

```

exit
queue 4
exit
queue 5
exit
queue 6
    rate cir 0 pir 4
exit
queue 7
    rate cir 3 pir 93
exit
queue 8
    rate cir 0 pir 3
exit
exit
network-queue "nq2-cbs" create
    queue 1
        rate cir 0 pir 32
        adaptation-rule cir max
    exit
    queue 2
    exit
    queue 3
    exit
    queue 4
    exit
    queue 5
    exit
    queue 6
        rate cir 0 pir 4
    exit
    queue 7
        rate cir 3 pir 93
    exit
    queue 8
        rate cir 0 pir 3
    exit
exit
exit
-----
*A:card-1>config>qos# info

```

8.5.3 Editing network queue QoS policies

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

8.6 Network queue QoS policy command reference

8.6.1 Command hierarchies

- [Configuration commands for 7210 SAS-T \(in network mode and access-uplink mode\)](#)
- [Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#)

- [Configuration commands for 7210 SAS-Sx 1/10GE and 7210 SAS-Sx 10/100GE](#)
- [Operational commands](#)
- [Show commands](#)

8.6.1.1 Configuration commands for 7210 SAS-T (in network mode and access-uplink mode)

```
config
- qos
- network-queue policy-name [create]
- description description-string
- no description
- queue queue-id
- adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
- no adaptation-rule
- rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
- no rate
```

8.6.1.2 Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

```
config
- qos
- network-queue policy-name [create]
- description description-string
- no description
- queue queue-id
- adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
- no adaptation-rule
- queue-mgmt name
- no queue-mgmt
- queue-mode queue-mode
- no queue-mode
- rate [cir cir-percent] [pir pir-percent]
- no rate
- weight weight
- no weight
```

8.6.1.3 Configuration commands for 7210 SAS-Sx 1/10GE and 7210 SAS-Sx 10/100GE

```
config
- qos
- network-queue policy-name [create]
- description description-string
- no description
- queue queue-id
- adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
- no adaptation-rule
- rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
- no rate
```

8.6.1.4 Operational commands

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

8.6.1.5 Show commands

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

8.6.2 Command descriptions

8.6.2.1 Configuration commands

8.6.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>network-queue

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

description-string

Specifies 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.

8.6.2.2 Network queue QoS policy commands

network-queue

Syntax

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

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure a network queue policy.

Network queue policies define the egress network queuing for the traffic egressing on the network ports and hybrid ports, and access-uplink ports (in access-uplink mode). Network queue policies define the bandwidth distribution for the various FC traffic egressing on the port. By default, network queue policy defines eight queues and a mapping of FC-to-queue.

Default

default

Parameters

policy-name

Specifies the name of the network queue policy. 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

create

Keyword to create a network queue policy.

8.6.2.3 Network queue QoS policy queue commands

queue

Syntax

queue *queue-id*

Context

config>qos>network-queue

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures a QoS network-queue policy queue.

On 7210 SAS-Mxp and 7210 SAS-T (both network and access-uplink mode), the behavior is as follows.

The queues are created by default (the user has no option to delete them) and the FCs are mapped to these queues as per [Table 31: Forwarding class to queue-ID map](#). Only one FC is mapped to one queue. Network queue carry both the unicast and multicast traffic and no segregation is done. The queues are scheduled are per the port scheduler policy associated with this port.

On the 7210 SAS-R6 and 7210 SAS-R12, the behavior is as follows.

The queues are created by default (the user has no option to delete them) and the FCs are mapped to queues, see [Table 31: Forwarding class to queue-ID map](#) for more information. Only one FC can be mapped to one queue. Queue ID 8 is the highest priority and queue ID 1 is the lowest priority. Network queues carry both the unicast and multicast traffic and no segregation is performed.

The hardware port scheduler prioritizes the queue according to the priority for each queue. High priority traffic should be mapped to high priority FCs. Mapping traffic to high priority FCs does not necessarily guarantee high priority treatment, because the scheduler policy can influence the relative priority among the queues. See [Schedulers on 7210 SAS-R6 and 7210 SAS-R12](#) for more information about scheduling behavior and the queue parameters considered by the scheduler.

On 7210 SAS-Sx 1/10GE: standalone and standalone-VC, the behavior is as follows.

The queues are created by default (the user has no option to delete them) and the FCs are mapped to these queues as per [Table 32: Forwarding class-to-queue ID map for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#). Network queue carry both the unicast and multicast traffic and a separate queue is used for them per FC. In other words, a total of 2 queues are allocated per FC, one queue each for unicast traffic and for multicast traffic. The queues are scheduled are per the port scheduler policy associated with this port.

Parameters

queue-id

Specifies the queue ID, 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 8

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]

no adaptation-rule

Context

config>qos>network-queue>queue

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command defines the method used by the system to derive the operational CIR and PIR rates when the queue is provisioned in hardware. For the **cir** and **pir** parameters, 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 CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **cir** and **pir** apply.

Default

adaptation-rule cir closest pir closest

Parameters

cir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced when adapting the CIR rate defined using the **queue queue-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used to derive the operational CIR rate for the queue. When the **cir** parameter is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

- Values**
- max** — Specifies that the operational CIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.
 - min** — Specifies that the operational CIR value is greater than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.
 - closest** — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

pir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced when adapting the PIR rate defined using the **queue-id rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR rate for the queue. When the **pir** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

- Values**
- max** — Specifies that the operational PIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.
 - min** — Specifies that the operational PIR value is greater than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.
 - closest** — Specifies that the operational PIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

queue-mgmt

Syntax

queue-mgmt *name*
no queue-mgmt

Context

config>qos>network-queue>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures a WRED policy for the specified queue.

The queue management policy is used to specify the queue buffer parameters and queue slope policy parameters.

The **no** form of this command associates the default SAP egress queue management policy with this queue.

Parameters

name

Specifies the name of the queue management policy, up to 32 characters.

queue-mode

Syntax

[no] **queue-mode** *queue-mode*

Context

config>qos>network-queue>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command determines whether the queue operates in strict or weighted mode. The **no** form of this command reverts the queue mode to the default value.

Default

weighted

Parameters

queue-mode

Specifies the mode of operation for the queue.

Values **strict** — If a queue is configured in strict mode, the scheduler schedules the queue in order of their priority in the 2 passes, the CIR loop and the PIR loop.

weighted — If a queue is configured in weighted mode, the scheduler examines these queues in two passes - CIR loop and a PIR loop. In the CIR loop, it distributes the available bandwidth to all the strict and then weighted queues in round-robin up to the configured CIR rate. It examines the weighted queues in the PIR loop, after examining all the strict queues and distributes the available bandwidth, if any, in the proportion of the configured weights.

rate

Syntax

rate [*cir cir-percent*] [*pir pir-percent*]

no rate

Context

config>qos>network-queue>queue

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the 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.

The **rate** command can be executed at any time, and alters the PIR and CIR rates for all queues created on the access ports.

The **no** form of this command reverts all queues created with the *queue-id* by association with the QoS policy to the default PIR (100) and CIR parameters (0).

Parameters

cir cir percent

Specifies the percentage of the guaranteed rate allowed for the queue. When the **rate** command is executed, a valid CIR setting must be explicitly defined. When the **rate** command has not been executed, the default CIR of 0 is assumed. Fractional values are not allowed and must be given as a positive integer.

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

Values 0 to 100

Default 0

pir pir percent

Specifies the percentage of the maximum rate allowed for the queue. When the **rate** command is executed, the PIR setting is optional. When the **rate** command has not been executed, or the PIR parameter is not explicitly specified, the default PIR of 100 is assumed. Fractional values are not allowed and must be given as a positive integer.

Values 1 to 100 percent

Default 100

weight

Syntax

weight *weight*

no rate

Context

config>qos>network-queue>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures the weight for the specified policy.

The configured weight determines the proportion of available bandwidth that is given to this queue in comparison to other queues contending for bandwidth at the same priority level.

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

Default

weight 1

Parameters

weight

Specifies the weight of the queue.

Values 1 to 15

8.6.2.4 Operational commands

copy

Syntax

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

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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 *src-name dst-name*

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

overwrite

Keyword 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, the following message is generated to indicate that the destination policy ID exists:

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

8.6.2.5 Show commands

network-queue

Syntax

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

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays network queue policy information.

Parameters

- network-queue-policy-name**
Specifies the name of the network queue policy. Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.
- detail**
Displays for each queue its rates, adaptation rule, and cbs details. It also shows FC to queue mapping details.

Output

The following outputs are examples of network queue policy information, and the associated tables describe the output fields:

- [Sample output 1, Table 59: Output fields: network queue labels](#)
- [Sample output 2, Table 60: Output fields: network queue policy](#)

Sample output 1

```
*A:card-1# show qos network-queue nql
=====
QoS Network Queue Policy
-----
Network Queue Policy (nql)
-----
Policy          : nql
Accounting      : packet-based
-----
Associations
-----
Port-id : 1/1/20
=====
*A:card-1#

*A:card-1# show qos network-queue nql detail
=====
QoS Network Queue Policy
=====
Network Queue Policy (nql)
-----
Policy          : nql
Accounting      : packet-based
Description     : this is a network-queue policy
-----
Queue  CIR      PIR      CBS
      CIR Rule  PIR Rule
-----
1      0         100      12.50
      closest  closest
2      0         100      12.50
      closest  closest
3      0         100      12.50
      closest  closest
4      0         100      12.50
      closest  closest
5      0         100      12.50
      closest  closest
```

```

6      0      100      12.50
7      closest closest
8      0      100      12.50
      closest closest
-----
FC      UCastQ
-----
be      1
l2      2
af      3
l1      4
h2      5
ef      6
h1      7
nc      8
-----
Associations
-----
Port-id : 1/1/20
=====
*A:card-1#
*A:card-1# show qos network-queue default detail
=====
QoS Network Queue Policy
-----
Network Queue Policy (default)
-----
Policy      : default

```

Table 59: Output fields: network queue labels

Label	Description
Policy	The policy name that uniquely identifies the policy
Description	A text string that helps identify the policy's context in the configuration file
Associations	Displays the physical port identifier where the network queue policy is applied
Queue	Displays the queue ID
CIR	Displays the committed information rate
PIR	Displays the peak information rate
CBS	Displays the committed burst size
FC	Displays FC to queue mapping

Sample output 2

```

*A:SAS>config>qos>network-queue# show qos network-queue default
=====
QoS Network Queue Policy
=====

```



```

-----
Network Queue Policy (default)
-----
Policy      : default
Accounting  : packet-based
Description : Default network queue QoS policy.
-----

Associations
-----
Port-id : 1/1/10
Port-id : 1/1/19
Port-id : 1/1/20
Port-id : 1/1/24
=====

*A:SAS>config>qos>network-queue# show qos network-queue default detail
=====
QoS Network Queue Policy
=====

Network Queue Policy (default)
-----
Policy      : default
Accounting  : packet-based
Description : Default network queue QoS policy.
-----

Queue Rates and Rules
-----

QueueId      CIR(%)      CIR Adpt Rule      PIR(%)      PIR Adpt Rule
-----
Queue1       0             closest            100         closest
Queue2       25            closest            100         closest
Queue3       25            closest            100         closest
Queue4       25            closest            100         closest
Queue5       100           closest            100         closest
Queue6       100           closest            100         closest
Queue7       10            closest            100         closest
Queue8       10            closest            10          closest
-----

Parent Details
-----

QueueId      Port      CIR Level      PIR Weight
-----
Queue1       True      1              1
Queue2       True      2              1
Queue3       True      3              1
Queue4       True      4              1
Queue5       True      5              1
Queue6       True      6              1
Queue7       True      7              1
Queue8       True      8              1
-----

High Slope
-----

QueueId      State      Start-Avg(%)      Max-Avg(%)      Max-Prob(%)
-----
Queue1       Down      70                 90               75
Queue2       Down      70                 90               75
Queue3       Down      70                 90               75
Queue4       Down      70                 90               75

```

Queue5	Down	70	90	75
Queue6	Down	70	90	75
Queue7	Down	70	90	75
Queue8	Down	70	90	75

Low Slope				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	50	75	75
Queue2	Down	50	75	75
Queue3	Down	50	75	75
Queue4	Down	50	75	75
Queue5	Down	50	75	75
Queue6	Down	50	75	75
Queue7	Down	50	75	75
Queue8	Down	50	75	75

Burst Sizes and Time Average Factor				

QueueId	CBS	MBS	Time Average Factor	

Queue1	def	def	7	
Queue2	def	def	7	
Queue3	def	def	7	
Queue4	def	def	7	
Queue5	def	def	7	
Queue6	def	def	7	
Queue7	def	def	7	
Queue8	def	def	7	

Table 60: Output fields: network queue policy

Label	Description
Policy	The policy name that uniquely identifies the policy
Description	A text string that helps identify the policy context in the configuration file
Associations	Displays the physical port identifier where the network queue policy is applied
Queue	Displays the queue ID
CIR(%)	Displays the committed information rate
CIR Adapt Rule	Displays the adaptation rule in use
PIR(%)	Displays the peak information rate
PIR Adapt Rule	Displays the adaptation rule in use
Port	Indicates if the parent scheduler is a port scheduler
CIR Level	Displays the priority of the queue in the CIR loop

Label	Description
PIR Weight	Displays the weight of the queue used in the PIR loop
High Slope	Displays the WRED high-slope parameters
Low Slope	Displays the WRED low-slope parameters
Burst Sizes (CBS/ MBS)	Displays the configured CBS and MBS value
Time Avg Factor	Displays the WRED Time Average Factor value in use
FC and UcastQ	Displays the FC-to-queue association

9 Service ingress QoS policies

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

9.1 Overview

There is one default service ingress policy. The default policy has two classification resources and one meter (the **num-qos-classifiers** set to value "2" and meter 1 is the default meter). The default policy uses CAM resources from the **ingress-internal-tcam>qos-sap-ingress-resource** pool for the classification of all traffic to the default FC, and rate-limits the traffic by using a policer that uses the default rate. SAP ingress policies with policing only are supported for SAPs configured on access ports and hybrid ports.



Note:

Queuing and shaping on SAP ingress is not supported on the 7210 SAS-T (access-uplink and network mode), 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.

Each policy can have up to 32 ingress meters. The default policies can be copied and modified but cannot be deleted. The default policies are identified as policy ID 1.

The default policies are applied to the service entry on creation (if applicable). For example, the default SAP ingress policy is applied to access ingress SAPs. You must explicitly associate other QoS policies.

9.1.1 Default SAP ingress policy

The default policy 1 maps all traffic to default forwarding class "be" and maps FC "be" to meter "1". Meter "1" is configured with **cir** 0 and **pir** max.

Output example

The following is a sample configuration output.

```
*A:7210-SAS>config>qos>sap-ingress# info detail
-----
description "Default SAP ingress QoS policy."
num-qos-classifiers 2
scope template
meter 1 create
    mode trtcm1
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs default
    cbs default
exit
default-fc "be"
```

9.1.2 SAP-ingress policy defaults

The following table lists the SAP-ingress policy defaults.

Table 61: SAP-ingress policy defaults

Field	Default
description	"Default SAP ingress QoS policy."
num-qos-classifiers	2
scope	template
meter	1
mode	trcm1
adaptation-rule	cir closest pir closest
rate	pir = max, cir= 0
mbs	default
cbs	default
default-fc	be

9.2 Resource allocation for SAP ingress policy

This sections describes resource allocation information for SAP ingress policies.

9.2.1 Use of index file by SAP QoS ingress policy

The 7210 SAS platforms use an index file to store the map that indicates the QoS resource allocation to SAPs. This file is used on reboot to ensure that all the SAPs that were created successfully before a reboot can be recreated after a reboot. Without an index file the system cannot ensure this (that is, without an index file it is possible that all the SAPs that were configured successfully may fail on a reboot after saving the configuration file). The file is stored in the flash memory.

On reboot, if the file is found, the system allocates resources in accordance with the stored map. If the file is not found, the system implements a best-fit algorithm and tries to allocate resources for all the SAPs on a first-come-first-served basis. If the file is not present, it is possible that the saved configuration will not execute successfully after the reboot, because the resources might not be allocated to all SAPs.



Note:

The index file used for the QoS map is different from the one used for storing interface indexes.

9.2.2 Use of the keyword “multipoint” for default meter “11”

The system allows sharing of a single meter for both unicast and multipoint traffic. Users can configure any of the available meters for multipoint traffic. The use of the **multipoint** keyword during meter creation has been deprecated, except for use with meter “11”, as described in the following paragraphs.

When the **multipoint** keyword is specified with meter “11” the system interprets it to be the default multipoint meter. The default multipoint meter is used for all FCs that do not have explicit multipoint meters configured. The system does the appropriate resource checks to ensure that resources needed to use multipoint meter with all the FCs are available before allowing this change.



Note:

- When **num-qos-classifiers** is set to a value of “2”, default multipoint meter “11” cannot be used because only a single meter is available for use.
- When associating a meter with an FC for broadcast, unknown-unicast, or multicast (BUM) traffic, the system does not validate whether the meter is a multipoint meter, thereby allowing users to use a single meter for unicast and BUM traffic. This implies efficient use of SAP ingress QoS resources. When the **multipoint** keyword is used, the system displays a warning to indicate that **multipoint** is an obsolete CLI command and is not saved in the configuration file, deprecating the use of the **multipoint** keyword with any meter other than the default.

9.2.2.1 Example uses of the multipoint meter

This section provides configuration examples of several uses of the multipoint meter.

Example

Example 1

```
*7210-SAS>config>qos# sap-ingress 12 create
*7210-SAS>config>qos>sap-ingress$ info
-----
      num-qos-classifiers 4
      meter 1 create
      exit
-----
*7210-SAS>config>qos>sap-ingress$
```

In example 1, all FCs in the SAP-ingress policy use the default meter 1 (for all traffic types). If the **configure qos sap-ingress id meter 11 multipoint** create command is executed, it attaches the default meter 11 with all the FCs defined in the SAP-ingress policy.

After configuring **meter 11 multipoint**, all the FCs in this policy use two meters: default meter 1, to meter unicast traffic for all the FCs; and meter 11, to meter BUM traffic for all the FCs. In this example, because only the default FC “be” is in use, the multipoint meter is used to meter BUM traffic associated with default FC “be”.

Example

The following example shows the policy after the configuration is changed.

```
*7210-SAS>config>qos# sap-ingress 12
*7210-SAS>config>qos>sap-ingress$ info
```

```
-----
        num-qos-classifiers 4
        meter 1 create
        exit
meter 11 multipoint create
-----
*7210-SAS>config>qos>sap-ingress$
```

Deleting the multipoint meter 11 removes all the FCs associated with the multicast-meter, assuming all the FCs are using the default multicast meter and do not have any other multicast meter explicitly configured. Executing the **configure qos sap-ingress id no meter 11** command disassociates meter 11 from the FCs, and the FCs use only meter 1 (if no other meter is configured explicitly).

Example

Example 2

```
*7210-SAS>config>qos# sap-ingress 12
*7210-SAS>config>qos>sap-ingress$ info
-----
configure> qos> sap-ingress 10 create
        meter 1 create
        exit
        meter 3 create
        exit
        default-fc be
        fc be
                meter 3
                multicast-meter 3
        exit
        fc af
                meter 3
        exit
exit
-----
```

Starting with the policy in example 2, if the user now executes the **configure qos sap-ingress id meter 11 multipoint create** command, the FC "be" continues to use meter 3 and the FC "af" uses meter 11 for BUM traffic. In this example, if the user executes the **configure qos sap-ingress id fc be no multicast-meter** command, the default meter 11 is also used for FC "be".

Example

Example 3

```
-----
configure> qos> sap-ingress 10 create
        meter 1 create
        exit
        meter 3 create
        exit

        default-fc be

        fc be
                meter 3
                unknown-meter 3
        exit
exit
-----
```

Upon executing the **configure qos sap-ingress id meter 11 multipoint create** command, FC "be" unknown-unicast traffic type continues to use meter 3, and broadcast and multicast traffic types use meter 11.

In example 3, if a broadcast-meter was initially configured in the SAP-ingress policy and was followed by executing the **configure qos sap-ingress id meter 11 multipoint create** command, FC "be" changes to use meter 11 for multicast traffic, and broadcast traffic continues to use meter 3 for unknown-unicast traffic and meter 3 for unicast traffic.

Also in example 3, if the user executes the **configure qos sap-ingress id fc be no unknown-meter** command, meter 3 is used for all traffic types classified to FC "be". But, if the default meter 11 is defined in the policy, FC "be" uses meter 11 for BUM traffic.

9.2.3 Service ingress meter selection rules

The following are rules for meter selection by different traffic types under various configurations.

9.2.3.1 Default policy

In the default policy, only meter "1" is defined. All FC and all traffic types use meter "1" by default. Meter "11" is not created by default and is not available for use.

Output example

The following is a sample configuration output.

```
*7210-SAS>config>qos# sap-ingress 1 create // Default policy
*7210-SAS>config>qos>sap-ingress$ info
-----
num-qos-classifiers 2
meter 1 create
exit
-----
*7210-SAS>config>qos>sap-ingress$
```

9.2.3.2 VPLS service without meter "11"

The following describes the usage of meters when meter "11" in a VPLS service is not configured in the policy:

- If an FC is created without explicit meters, the default meter "1" is used for unicast traffic and for multipoint traffic types (such as broadcast, multicast and unknown-unicast traffic).
- If an FC is created with an explicit unicast meter, that meter is used for unicast traffic and for multipoint traffic types (such as broadcast, multicast and unknown-unicast traffic).
- If an FC is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively and use the unicast meter for all other traffic types.
- If an FC is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other traffic types.
- If an FC is created with an explicit unicast meter, an explicit broadcast meter, and an explicit multicast meter, use these meters for unicast, broadcast and multicast traffic types respectively. Unknown unicast traffic type will use the explicitly defined multicast meter.

- If an FC is created with an explicit unicast meter, an explicit broadcast meter, an explicit unknown-unicast meter, and an explicit multicast meter, use these meters for unicast, broadcast, unknown-unicast and multicast traffic types respectively.

9.2.3.3 VPLS service with meter "11"

The following describes the usage of meters when meter "11" in a VPLS service is defined in the policy:

- If an FC is created without explicit meters, use the default meter "1" for unicast traffic and default meter "11" for all other traffic types (such as broadcast, multicast and unknown-unicast).
- If an FC is created with an explicit unicast meter, use that meter for unicast traffic and use default meter "11" for all other traffic types.
- If an FC is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively and use meter "11" for all other traffic types.
- If an FC is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic.
- If an FC is created with an explicit unicast meter, an explicit broadcast meter, and an explicit multicast meter, use these meters for unicast, broadcast and multicast traffic types respectively. Unknown unicast traffic type will use the explicitly defined multicast meter.
- If an FC is created with an explicit unicast meter, an explicit broadcast meter, an explicit unknown-unicast meter, and an explicit multicast meter, use these meters for unicast, broadcast, unknown-unicast and multicast traffic types respectively.

9.2.3.4 Epipe, IES, and VPRN services without PIM

The following are rules for meter selection for Epipe, IES and VPRN services:

- A multipoint meter cannot be used. A multipoint meter configured in a policy is not used when the policy is applied to a SAP in an Epipe service.
- All FCs associated with a meter always use the unicast meter.



Note:

These rules apply to IES services when PIM is not enabled in the service.

9.2.3.5 IES and VPRN services with PIM/multicast and without meter "11"

The following are rules for meter selection for IES and VPRN services when PIM/multicast is enabled in the service and describes the usage of meters when meter "11" is not configured in the policy:

- If an FC is created without explicit meters, the default meter "1" is used for unicast traffic and multicast traffic.
- If an FC is created with an explicit unicast meter, that meter is used for unicast traffic and for multicast traffic.
- If an FC is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for multicast traffic.

9.2.3.6 IES and VPRN services with PIM/multicast and meter "11"

The following are rules for meter selection for IES and VPRN services when PIM/multicast is enabled in the service and when meter "11" is defined in the policy:

- If an FC is created without explicit meters, use the default meter "1" for unicast traffic and default meter "11" for multicast traffic.
- If an FC is created with an explicit unicast meter, use that meter for unicast traffic and use default meter "11" for multicast traffic.
- If an FC is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for multicast traffic.

9.2.4 Service ingress policy configuration considerations

The **num-qos-classifiers** parameter cannot be modified when the policy is in use (for example, when it is associated with a SAP). Other parameters in the SAP ingress policy can be changed.

When changing other parameters for a policy that is in use (for example, fc meter map or fc classification match criteria entries), the system recomputes the resources required to accommodate the change. If the resources required exceeds the configured value for **num-qos-classifiers**, the change is not allowed.

If more resources are needed than the value configured in **num-qos-classifiers** for a existing policy, then the following options are available:

- Copy the existing policy to a new policy, modify the **num-qos-classifiers** parameter, modify the match criteria entries, and modify the SAP configuration to associate it with the new policy.
- Ensure the existing policy is not in use by any SAP (if required, change the SAP configuration to disable the use of the QoS policy with the **no qos** form of this command), change all the required parameters, and finally modify the SAP configuration to use the policy again.



Note:

Both options above have side effects; for example, they can reset the statistics associated with the meters and can potentially cause existing traffic classification not to take effect. However, the system ensures that the default policy is used during the period when policy changes are being made after the two options are performed.

The following items are additional service ingress policy configuration considerations:

- In releases before Release 3.0R1, the system always computes the number of resources (like classifiers and meters) required by a policy assuming the number of resources will be used in a VPLS service. This allows the policy to be applied to either an Epipe or VPLS service.
- From Release 3.0R1 onwards, on creation of SAP ingress policy, the system does not compute the number of resources required by a policy and validate the number against resources available in the system. The system validates the resources needed only when the SAP ingress policy is attached to a SAP. If enough resources are available the association succeeds, otherwise the system fails the CLI command. Based on the service in which the SAP is configured (such as VLL, VPLS, and so on), for the same SAP ingress policy the amount of resources required is different. The system validates that the amount of QoS resources specified in the **num-qos-classifiers** command is sufficient for the match criteria, forwarding class, and service specified and that the resources are available in hardware. On failure of the validation, the system disallows the association of the SAP ingress policy with the SAP.

- The match criteria type (that is, mac-, ipv4-, and ipv6-criteria) cannot be changed when the SAP ingress QoS policy is in use. For example, if the match-criteria is set to ipv4-criteria and the policy is associated with a SAP, then the ipv6-criteria or mac-criteria cannot be enabled in the same policy. If there is a need to change the criteria, users must remove the association and then change the SAP ingress policy to use the new match criteria. For SAPs configured in VPRN services, the computation of resources is similar to an SAP configured in an Epipe service.

See [Resource allocation for service ingress QoS policies using CAM-based classification](#) for more information.

9.2.5 Resource allocation for service ingress QoS policies using CAM-based classification

The available global pool of ingress internal CAM hardware resources can be allocated based on user needs and shared among different features, such as SAP ingress QoS policy, ingress ACLs, and so on. Use the **configure>system>resource-profile>ingress-internal-tcam>qos-sap-ingress-resource** command to allocate SAP ingress QoS classification and meter resources from this global pool.

In addition, resources can be allocated for different SAP ingress QoS policy classification match criteria based on operator needs. Users can modify resource allocation to scale the number of entries available per match criteria or to scale the number of SAPs. Resources from the global ingress internal CAM pool are allocated in fixed slices. The number of classification entries and meters per slice varies across 7210 SAS platforms.

For more information about number of slices that can be allocated to a feature, see the "System Resource Allocation" section in the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide*. For example, on a 7210 SAS-R6 IMM-b card, each slice allows 256 classification entries and 128 meters. The number of slices allotted for a SAP ingress QoS policy is configured using the **configure>system>resource-profile>ingress-internal-tcam>qos-sap-ingress-resource** command.

Users can configure the resources available for SAP ingress QoS policies and can limit the amount of resources used per match criteria supported for SAP ingress QoS policies. A specific slice can be used for either MAC criteria or IP criteria or IPv6 criteria, or both MAC and IP criteria. Allocation of classification entries also allocates meter/policer resources that are used to implement per FC per traffic type policing.



Note:

For SAP ingress classification, in addition to CAM-based resource allocation, the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 support table-based resource allocation for DSCP-classification on SAP ingress. See [Table-based classification using dot1p and IP DSCP for assigning FC and profile on SAP ingress for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information.

By default, the system allocates resources for SAP ingress QoS policies to maintain backward compatibility with Release 4.0 and allocates resources for MAC criteria and IP criteria (by setting *num-resources* to **max**). Setting the *num-resources* value to **max** allows each match criteria to use the available SAP ingress QoS resources on a first-come-first-served model. By default, system does not allocate resources for use by ingress IPv6 classification or for using both IP (any) and MAC (any) criteria in a single SAP ingress policy. You must allocate resources before associating a SAP with an IPv6 SAP ingress policy or both IP (any) and MAC (any) criteria in a SAP ingress policy. Until appropriate resources are allocated, attempts to associate the policy with the SAP will fail.

If the **configure>system>resource-profile>qos-sap-ingress-resource** command is used to allocate resources for SAP ingress QoS policies, the system allocates resources in slices for a fixed number of

entries; the entries allocated per slice is platform-dependent. The use of these entries by different types of match criteria is described in the following table.

Table 62: SAP ingress resource allocation and match criteria types

Type of match criteria	Description
mac-criteria (any)	<p>Before using SAP-ingress policies with mac-criteria (any), use the configure system resource-profile ingress-internal-tcam qos-sap-ingress-resource mac-match-enable CLI command to allocate resources from the SAP-ingress QoS resource pool. Resources are allocated for SAP-ingress policies that use mac-criteria (any and dot1p-only). Each entry in the SAP-ingress QoS policy that is configured to use mac-criteria uses one entry from the slices in the hardware resource pool allocated to mac-criteria.</p> <p>For example, assume a SAP ingress QoS policy is configured to use mac-criteria with 50 entries and uses the configure> system> resource-profile> ingress-internal-tcam> qos-sap-ingress-resource> mac-match-enable 1 CLI command to configure one slice for use by mac-criteria (allowing a total of 512 entries for use by policies that use mac-criteria). In this case, the user can configure 10 SAPs that use the mac-criteria SAP ingress policy and consume 500 entries.</p>
ipv4-criteria (any)	<p>The use of ipv4-criteria (any) match criteria is the same as the mac-criteria. Use the configure system resource-profile ingress-internal-tcam qos-sap-ingress-resource ipv4-match-enable CLI command to allocate resources. Additionally, IPv4 criteria can share entries allocated for IPv6 criteria. The SR OS automatically allocates entries from an IPv6 criteria slice to IPv4 criteria policies, if no entries are available in the allocated IPv4 criteria slices and no slices are available for allocation to IPv4 criteria from the SAP-ingress QoS resource pool. If an IPv4 criteria entry uses IPv6 criteria slices, the number of hardware entries used is the same as required by an IPv6 criteria entry (see ipv6-criteria (any) for more information).</p>
ipv6-criteria (any)	<p>Before using the ipv6-criteria match criteria, use the configure> system> resource-profile> ingress-internal-tcam> qos-sap-ingress-resource> ipv6-ipv4-match-enable CLI command, and specify the ipv6 keyword for the num-qos-classifiers command to allocate resources from the SAP ingress QoS resource pool.</p> <p>Each ipv4-criteria match entry or ipv6-criteria configured in the QoS policy that uses ipv6-criteria uses two (2) entries from the slices allocated for use by ipv6-criteria (128-bit) in the hardware. The system allocates entries from the ipv6-criteria pool in the following cases:</p> <ul style="list-style-type: none"> • if the SAP ingress QoS policy uses ipv6-criteria entries (any or IPv6 DSCP) and ipv4-criteria entries (any or IPv4 DSCP) • if the SAP ingress QoS policy uses only ipv6-criteria (any or IPv6 DSCP) • if the SAP ingress QoS policy uses ipv4-criteria (any) and there are no resources available in the ipv4-criteria (see ipv4-criteria (any) for more information) <p>For example, assume a QoS policy is configured to use ipv6-criteria with 50 entries and the configure>system>resource-profile>ingress-internal-tcam>qos-sap-ingress-resource>ipv4-ipv6-128-match-enable 1 command is used to configure one slice for ipv6-criteria. This allows a total of 256 entries for use by SAPs using SAP ingress QoS policies with ipv6-criteria (because each IPv6 entry uses 2 entries in hardware). In this example, the user can configure five (5) SAPs that use this policy</p>

Type of match criteria	Description
	and consume a total of 250 entries. These resources can be shared with policies that use IPv4 criteria, though each IPv4 criteria entry consumes two (2) entries in the hardware. IPv4 criteria policies can consume spare IPv6 resources; however, if a larger number of IPv4 criteria policies are planned, it is good practice to allocate more resources for use with IPv4 criteria. ¹⁶
IPv4 (any) and MAC (any) match	<p>Before using IP-criteria (any) and MAC-criteria (any) in a single policy, use the configure>system>resource-profile>ingress-internal-tcam>qos-sap-ingress-resource>ipv4-mac-match-enable command to allocate resources from the SAP ingress QoS resource pool. Each ipv4-criteria match entry or MAC-criteria match entry configured in the QoS policy uses two (2) entries from the allocated slices. The system allocates entries from the ipv4-mac-match-enable pool if the SAP ingress QoS policy uses both ip-criteria (any) and ipv6-criteria (any).</p> <p>The system also allocates entries for all other criteria if there are no resources available to use in the pool allocated to those criteria. That is, if no resources are available in other pools, the following criteria are allocated resources from this pool: only mac-criteria any, only ip-criteria any, mac-criteria dot1p-only, ip-criteria dscp-only, ipv4-criteria dscp-only. ¹⁷</p>
dot1p-only, IPv4 dscp-only, IPv6 dscp-only, and default SAP ingress QoS policies	<p>Use the dot1p-only or dscp-only option if only dot1p bits or only IP DSCP bits or only IP precedence bits are used for SAP ingress classification. This facilitates efficient use of available hardware resources and better scaling. SAP ingress policies that use only dot1p bits or only IPv4/IPv6 DSCP or IPv4/IPv6 precedence or default SAP ingress QoS policy bits can use the resources from slices currently allocated for use by either IP-criteria or MAC-criteria or IPv6 criteria. The following is a list of special cases for resource allocation for default, dot1p-only, and dscp-only SAP ingress policies are as follows:</p> <ul style="list-style-type: none"> • If no slices are available to accommodate a SAP associated with default or dot1p-only or a dscp-only SAP ingress policy, the SR OS allocates resources against mac-criteria if the SAP is configured in a VLL or VPLS service. The SR OS uses the required number of entries for this policy. The remaining entries are available for SAPs that use mac-criteria or that use only dot1p or only IPv4/IPv6 DSCP or that use a default policy. • If no slices are available to accommodate a SAP associated with default, dot1p-only, or a dscp-only SAP ingress policy, the SR OS allocates resources against ipv4-criteria if the SAP is configured in an IES, RVPLS, or a VPRN service. The SR OS uses the required number of entries for this policy. The remaining entries are available for SAPs that use ipv4-criteria or that use only IPv4/IPv6 DSCP or only dot1p criteria or that use default policy. • For a SAP ingress QoS policy, if the enable-table-classification command is set to use table-based classification and TCAM-based meters, and there are no resources available in the pool of slices allocated for ip-dscp-port-if-match criteria, a new slice

¹⁶ When a slice is allocated for IPv6 criteria, the system automatically adjusts the number of available classification entries in that slice to half the number of total entries available. For example, on the 7210 SAS-T, if a slice with a total of 256 entries is allocated for IPv6 criteria, the actual number of classification entries available for IPv6 classification is only 128. The number of meters available is 128 in this example.

¹⁷ Resources are not allocated to ipv6-criteria (any) from this pool.

Type of match criteria	Description
	is allocated and set to the ip-dscp-port-if-match-enable criteria. The SR OS uses the required number of entries from this slice for the policy being configured. The remaining entries are used for SAPs configured with default sap ingress QoS policy, or for SAPs with the enable-table-classification command set, which uses dot1p or IP DSCP table-based classification with TCAM meters.

The SAP ingress resource slices described in this section are different from the resources that are configured using the **num-qos-classifiers** command. The **num-qos-classifiers** command sets the limit on the resources needed per SAP ingress QoS policy. The **qos-sap-ingress-resource** resources set the maximum limit on the resources available to all the SAP ingress policies that are in use simultaneously on the system.

The SR OS manages the resource slices allocated to the SAP ingress QoS policy pool and allocates the slice entries when a SAP ingress QoS policy is associated with a SAP. In other words, a SAP specifies the number of QoS resources it needs using the **num-qos-classifiers** command (in the SAP ingress policy), while the system allocates the resources required by the SAP from the **qos-sap-ingress-resource** slices, depending on whether the SAP ingress policy uses ip-criteria or mac-criteria or ipv6-criteria.

Use the **tools dump system-resources** command to view the current usage and availability of system resources. One or more entries per slice are reserved for system use.

9.2.5.1 Resource configuration guidelines for service ingress QoS policies using CAM-based classification

When a SAP is associated with a default SAP ingress QoS policy and there are no resources available in the pool of slices already allocated for different criteria that are in use, a new slice is allocated and set to either mac-match criteria, ipv4-match criteria, or ip-dscp-port-if-match criteria. This allocation can result in a single slice getting consumed and becoming unavailable for other classification criteria even if the mac-match criteria, ipv4-match criteria, or ip-dscp-port-if-match criteria are not used. To prevent this scenario, the SAP ingress resource configuration can be set to the specific number of slices for each criteria in use so that the SR OS can allocate the slices based on user requirement without allocating resources for any of the mac-match criteria, ipv4-match criteria, or ip-dscp-port-if-match criteria.

9.2.6 Computation of resources used per SAP ingress policy for CAM-based classification

Users can configure the number of classification entries the SAP requires (TQ). The value of TQ is set using the **num-qos-classifiers** command, where TQ is the total number of QoS resources required by the SAP. To determine TQ, see [Determining the number of policers/meters per policy \(TP\)](#).

Number of meters allocated automatically by system = $TQ / 2$ (up to a maximum of 32 meters)

To calculate the number of SAPs allowed, assume all SAPs are configured to use "TQ" QoS resources per SAP.

Number of SAPs allowed = maximum classification entries / TQ



Note:

The number of SAPs calculated using the equation above is subject to system limits. The above equation is used to derive the limit on the number of SAPs due to QoS resources only.

Users can mix and match SAPs with different QoS resources (that is, use different values of TQ).

The following criteria determine the number of QoS resources allocated for a SAP:

- number of match-criteria entries used to identify the FC
- number of FCs to use and number of traffic types to be policed per FC
- amount of hardware classification resources needed per entry configured by the user; see [Resource allocation for service ingress QoS policies using CAM-based classification](#) for more information about resources needed per match entry, which varies based on the different match criteria in use

Only FCs that are in use by the match-criteria classification entries are considered for the number of FCs. These FCs are referred to as "FC in use".

9.2.6.1 Determining the number of classification entries

This section describes the rules and methods of determining the number of classification entries.

9.2.6.1.1 Rules for a SAP in a VPLS

Knowing the number of traffic types to use per "FC in use", apply the following rules for a SAP in a VPLS service to determine the number of classification entries per "FC in use":

- If an FC is in use and is created without explicit meters, use default meter "1" for unicast traffic and for all other traffic types (that is, broadcast, multicast and unknown-unicast). This requires one classification entry in hardware. This assumes default multipoint meter "11" is not created by the user.
- If an FC is in use and is created without explicit meters, use default meter "1" for unicast traffic and default meter "11" (assuming meter "11" is created by the user), for all other traffic types (that is, broadcast, multicast and unknown-unicast). This requires two classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and for all other traffic types (that is, broadcast, multicast and unknown-unicast). This requires one classification entries in hardware. This assumes default multipoint meter "11" is not created by the user.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter "11" (assuming meter "11" is created by the user) for all other traffic types. This requires two classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively and use the unicast meter for all other traffic types (that is, multicast and unknown-unicast). This requires two classification entries in hardware. This assumes that the default multipoint meter "11" is not created by the user.
- If an FC is in use and is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively and use meter "11" (assuming meter "11" is created by the user) for all other traffic types. This requires three classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic. This requires two classification entries in hardware.

- If an FC is in use and is created with an explicit unicast meter, an explicit broadcast meter, and an explicit multicast meter, use these meters for unicast, broadcast and multicast traffic types respectively. Unknown unicast traffic type will use the explicitly defined multicast meter. This requires three classification entries in hardware.

9.2.6.1.2 Rules for a SAP in a VLL, VPRN, or IES service with PIM disabled

Apply the following rules for a SAP in a VLL, VPRN, or IES service with PIM disabled, to determine the number of classification entries per FC:

- Multipoint meters cannot be used. Multipoint meter configured in a policy is not used when the policy is applied to a SAP in an Epipe service.
- All FCs in use and associated with a meter always use the unicast meter. Therefore, all FCs in use utilize only one classification entry in the hardware.

9.2.6.1.3 Rules for a SAP in an IES or VPRN service with PIM/multicast enabled

Knowing the number of traffic types to use per "FC in use", apply the following rules for a SAP in an IES or VPRN service enabled with PIM/multicast enabled to determine the number of classification entries per FC in use:

- If an FC is in use and is created without explicit meters, use default meter "1" for unicast traffic and for multicast traffic. This requires one classification entry in hardware. This assumes default multipoint meter "11" is not created by the user.
- If an FC is in use and is created without explicit meters, use default meter "1" for unicast traffic and default meter "11" (assuming meter "11" is created by the user), for multicast traffic. This requires two classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and for multicast traffic. This requires one classification entries in hardware. This assumes default multipoint meter "11" is not created by the user.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter "11" (assuming meter "11" is created by the user) for multicast traffic. This requires two classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for multicast traffic. This requires two classification entries in hardware.

9.2.6.1.4 Calculating the number of classification entries per FC

Apply the rules (above) to determine the number of classification entries per FC (C_i) (only for the "FCs in use") using the following equation:

$$C(i) = \sum FC_i (\text{unicast}) + FC_i (\text{multicast}) + FC_i (\text{broadcast}) + FC_i (\text{unknown_unicast})$$

$i = nc, h1, ef, h2, l1, af, l2, be$

where:

- FC_i (unicast), FC_i (multicast), FC_i (broadcast), and FC_i (unknown-unicast) are set to a value of 1 if the FC uses a classifier to identify traffic-type unicast, multicast, broadcast, and unknown-unicast, respectively
- FC_i (unicast), FC_i (multicast), FC_i (broadcast), and FC_i (unknown-unicast) are set to a value of 0 if the FC does not use a classifier to identify the traffic-type

If the user does not configure meters explicitly for the FC and multipoint meter "11" is not created, the default unicast meter is used for all traffic types and therefore, only one classification entry in hardware is required by the FC. If the user does not configure meters explicitly for the FC and multicast meter "11" is created, the default unicast meter and the default multicast meter are used. Therefore, by default, two classification entries in hardware are required by an FC.

9.2.6.1.5 Determining the number of classification entries per policy (TC)

Taking into account the number of match criteria and the number of FCs used, use the equation below to determine the total number of classification entries per policy, for example:

$$TC = \sum E(i) * C(i)$$

$i = nc, h1, ef, h2, l1, af, l2, be$

where:

- $E(i)$ is the number of match-criteria entries that classify packets to FC_i . For 7210 SAS platforms, the maximum number of classification entries per policy can be 64 (including default).
- $C(i)$ is the number of classification entries that are required by FC_i to identify different traffic types.

9.2.6.1.6 Determining the number of policers/meters per policy (TP)

Determine the number of policers or meters to use (TP). A maximum of 32 meters per policy are available. The number of policers/meters is determined by the number of meters associated with FCs in the SAP-ingress QoS policy.

Use the values of TC and TP to determine the required number of QoS resources (TQ).

Only those meters associated with FCs are considered for the number of meters. Note that only "FCs in use" are considered.

$$\text{Total QoS resources required (TQ)} = \max [(TC), (2 * TP)]$$

The resulting number is rounded up to the next multiple of "2" greater than TQ obtained above. For example, if $TC = 5$ and $TP = 2$, then $\max (5, (2 * 2))$ is 5, and TQ is rounded up to 6.

The user configures TQ value using **num-qos-classifiers** command.

For more information and examples about resource calculation, see the following sections:

- [Service ingress policy configuration considerations](#)
- [Examples: calculating resources required for CAM-based classification](#)
- [Examples: calculating resources required for IP DSCP table-based classification with CAM-based policing \(7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12\).](#)

9.3 Table-based classification using dot1p and IP DSCP for assigning FC and profile on SAP ingress for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Using an IP DSCP classification policy and a dot1p classification policy is another method used to assign the FC and profile on SAP ingress for use with color-aware meters. Similar to network ingress, you can use an IP DSCP classification policy and a dot1p classification policy to assign the FC and profile on SAP ingress for use with color-aware meters. In this section, this classification is also called table-based classification. Table-based classification is supported on the 7210 SAS-Mxp both in low-sap-scale mode and high-sap-scale mode. Table-based classification and CAM-based classification are mutually exclusive.

This is supported along with the capability to use other header fields in MAC and IP packets through the use of the **mac-criteria**, **ipv4-criteria**, and **ipv6-criteria** commands. When using SAP-ingress color-aware meters and policers, users can configure DEI to assign the initial profile on ingress and can configure either MAC criteria or IP criteria (or both) to assign the FC.

To configure IP DSCP classification, users create a DSCP classification policy, associate the policy with a SAP ingress QoS policy or an Ethernet ingress port or L3 interface (as applicable), apply the SAP ingress QoS policy to the SAP or port, and enable the use of the policy.

Similarly, to configure dot1p classification, users create a dot1p classification policy, associate the policy with a SAP ingress QoS policy or an Ethernet ingress port, apply the SAP ingress QoS policy to the SAP or port, and enable the use of the policy.

DSCP and dot1p classification use classification resources from a table classifier and, potentially, a lesser number of resources from the CAM resources, thereby saving CAM resources for other purposes. When a table-based classification policy is enabled, CAM-based classifications from the SAP ingress QoS policy are ignored.

The following topics describe the support for FC and profile assignment based on IP DSCP and dot1p on SAP ingress.

9.3.1 IP DSCP and dot1p classification policy support

This section describes IP DSCP and dot1p classification support.

For SAPs in Layer 2 services (Epipe and VPLS) configured to use table-based classification, you can use the **config>qos>sap-ingress>table-classification-criteria** CLI command, which provides the option to select one of the following: dot1p or IP DSCP, or both dot1p and IP DSCP, or none. This command is applicable only to SAPs configured in Layer 2 services; it is ignored for Layer 3 services (IES, VPRN and RVPLS services).

The following behavior is supported with table-based classification for SAPs configured in Layer 2 services depending on the configuration option selected with the **config>qos>sap-ingress>table-classification-criteria** CLI command:

- If **none** is configured, use **default-fc fc-name profile out** (from the SAP ingress policy).
- If **use-dscp** is configured, use the following policies:
 - DSCP classification policy for IP packets
 - **default-fc fc-name profile out** (from the SAP ingress policy) for non-IP packets
- If **use-dot1p** is configured, use the following policies:

- dot1p classification policy for all tagged packets (IP and non-IP)
- **default-fc** *fc-name* **profile out** (from the SAP ingress policy) for untagged packets
- If **both-dscp-dot1p** is configured, use the following policies:
 - DSCP classification policy for IP packets
 - dot1p classification policy for non-IP tagged packets
 - **default-fc** *fc-name* **profile out** (from the SAP ingress policy) for non-IP untagged traffic

Table-based classification for SAPs configured in Layer 3 services (IES and VPRN) does not allow classification options with the **config>qos>sap-ingress>table-classification-criteria** CLI command. It must always be set to both-dscp-dot1p. The following behavior is supported:

- IP DSCP-based classification to assign FC and profile to an IP packet (both IPv4 and IPv6) received on SAP ingress



Note:

All SAPs in an RVPLS service can use either table-based classification or CAM-based classification entries. There cannot be a mix of SAPs where some SAPs use table-based classification and others use CAM entries.

The following is the order of match for packets with table-based classification in Layer-3 services:

1. IP packets are matched against IP DSCP entries. Ideally, packets match one of the explicitly configured user entries that classifies packets to the configured FC and profile. If there is no match, packets are assigned a default-dscp-fc and profile. If a packet is a non-IP packet, go to 2.
2. Non-IP tagged Ethernet packets that do not match any IP DSCP entries are matched against dot1p values. Ideally, packets match one of the explicitly configured user entries that classifies packets to the configured FC and profile. If there is no match, packets are assigned a default-dot1p-fc and profile. Alternatively, users have an option to use Drop Eligible Indicator (DEI) to assign the profile. If a packet is an untagged non-IP packet, go to item 3.
3. All non-IP untagged Ethernet packets are assigned a **default-fc/untagged-fc** (user-configurable) value for FC and profile out.

9.3.1.1 Default-FC assignment rules for SAPs in Layer 3 services

The following **default-fc** assignment rules apply for SAPs configured in Layer 3 services, including RVPLS services:

- Bridged IP packets processed in an RVPLS context that do not match any of the explicitly configured DSCP classification entries in the access port DSCP classification policy are assigned the **default-dscp-fc** configured in the policy. The profile assigned on ingress is in accordance with the user configuration (for example, **default-dscp-fc "be" profile out**).
- Bridged, tagged non-IP Layer 2 Ethernet packets, which are processed in an RVPLS context and that do not match any of the explicitly configured dot1p classification entries in the access port dot1p classification policy, are assigned the **default-dot1p-fc** configured in the policy. The profile assigned on ingress is in accordance with the user configuration (for example, **default-dot1p-fc "be" profile out**). Bridged, untagged non-IP Layer 2 Ethernet packets, which are processed in an RVPLS context, will not match any of the explicitly configured dot1p classification entries in the access port dot1p classification policy. They are assigned access port untagged-fc profile.

- Routed IP packets processed in an RVPLS context that do not match any of the explicitly configured DSCP classification entries in the DSCP classification policy associated with the interface using the **routed-override-qos-policy** command are assigned the **default-dscp-fc** configured in the policy. The profile assigned on ingress is in accordance with user configuration in the DSCP classification policy.
- IP packets processed in an IES or VPRN service context that do not match any of the explicitly configured DSCP classification entries in the DSCP classification policy associated with SAP ingress policy are assigned the **default-dscp-fc** configured in the DSCP classification policy. The profile assigned on ingress is in accordance with user configuration in the DSCP classification policy.

9.3.2 Precedence rules for DEI assignments on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

The following precedence rules pertain to the use of DSCP/dot1p classification policies and DEI profile assignment:

- Assigning a profile using an IP DSCP/dot1p classification policy has precedence over assigning a profile using DEI bits. That is, if both a DEI-based initial profile assignment and a DSCP-based initial profile assignment are enabled on a port, the profile assigned by DSCP/dot1p-classification policy takes precedence for IP and non-IP packets, respectively.
- The CLI command to enable DEI-based profile assignment (**configure>port> ethernet>enable-dei**) can be used to assign a profile for SAPs configured to use CAM-based classification entries:
 - A port can have a mix of SAPs with some SAPs configured to use DEI for profile assignment with CAM-based classification and other SAPs configured to use table-based classification for both FC and profile assignment.
 - The **enable-dei** command does not affect SAPs that use table-based classification. That is, a profile assignment for SAPs configured to use table-based classification always uses a DSCP or dot1p classification policy.

9.3.3 Creating an IP DSCP and dot1p classification policy

IP DSCP and dot1p classification policies can be used at SAP ingress, allowing users to define the mapping of an IP DSCP or dot1p value to an FC and profile.

The default values for the **default-dscp-fc** and **default-dot1p-fc** command are FC "be" and profile "out". Newly-created classification policies contain the default "be" and "out" values as the default entries. The **default-dscp-fc** or **default-dot1p-fc** command assigns the default FC to any IP DSCP or dot1p value that is not explicitly configured by a user.

Up to 50 unique DSCP or dot1p classification policies can be supported on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

Use the following syntax to create a DSCP classification QoS policy.

```
configure> qos>
  dscp-classification classification-id [create]
    dscp dscp-name [fc fc-name] [profile [in | out]]
    default-dscp-fc fc profile [in |out]
```

The following example shows the command usage to create a DSCP classification QoS policy.

Example:

```
configure> qos>
  dscp-classification 101 create
  dscp af42 fc af profile in
  dscp af43 fc af profile out
  dscp af32 fc h1 profile in
  dscp af33 fc h1 profile out
  dscp ncl fc nc profile in
  default-dscp-fc be profile out
exit
```

Use the following syntax to create a dot1p classification QoS policy.

```
configure> qos>
  dot1p-classification classification-id [create]
  dot1p dot1p-name [fc fc-name] [profile [in | out]
  default-dot1p-fc fc profile [in |out]
```

Example:

The following example shows the command usage to create a dot1p classification policy.

```
configure> qos>
  dot1p-classification 101 create
  dot1p af42 fc af profile in
  dot1p af43 fc af profile out
  dot1p af32 fc h1 profile in
  dot1p af33 fc h1 profile out
  dot1p ncl fc nc profile in
  default-dot1p-fc be profile out
exit
```

9.3.4 CAM resource usage for IP DSCP and dot1p classification policies on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

The resources for IP DSCP and dot1p classification policies are taken from hardware tables, which is referred to as table-based classification to differentiate it from CAM-based classification. Table-based resources do not use many CAM entries for classification. Only a fixed number of CAM resources is needed to match the FC and traffic type and to assign a meter/policer. Table-based classification uses CAM resources more efficiently than CAM-based classification:

- Users can enable an IP DSCP and dot1p classification policy on a per-SAP or per-port basis. See [Associating a DSCP or dot1p classification policy](#).
- The rules for associating an IP DSCP and dot1p classification policy with SAPs for different services are different. See [Assigning and enabling DSCP and dot1p classification policies to a SAP](#).

To calculate the number of resources needed, see the following sections:

- [Computation of resources used per SAP ingress policy for CAM-based classification](#)
- [Examples: calculating resources required for CAM-based classification](#)
- [Examples: calculating resources required for IP DSCP table-based classification with CAM-based policing \(7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12\)](#)

9.3.4.1 CAM resource allocation for table-based classification

When table-based classification is used from the SAP ingress resource pool, CAM resources are used to match the FC and the traffic type (unicast, broadcast, multicast, and unknown) and to assign a meter to the FC. The maximum number of DSCP entries required is 36 entries [8 FCs * 4 traffic types + 1 default FC * 4 traffic types). The maximum number of CAM resources required is 32, which assumes the use of VPLS service with one meter per traffic type, regardless of the number of IP DSCP classification entries. In other words, even if 64 IP DSCP values are matched, only 32 classification resources in the SAP ingress CAM resource pool are needed. For Epipe services, the number of entries is reduced to nine (8 FCs + 1 default FC) since all traffic is treated as unicast.

To allocate resources to meters for a SAP ingress QoS policy that is using table-based DSCP classification, use the **ip-dscp-port-if-match-enable** *num-resources* command, found under the **configure> system> resource-profile> ingress-internal-tcam> qos-sap-ingress-resource** context. The command supports up to 10 resources (meters).

Use the following syntax to allocate resources to meters for table-based classification.

```
configure> system> resource-profile>
    ingress-internal-tcam>
        qos-sap-ingress-resource>
            ip-dscp-port-if-match-enable num-resources
```

The following example shows the command usage to allocate resources to meters for table-based classification.

Example:

```
configure> system> resource-profile>
    ingress-internal-tcam>
        qos-sap-ingress-resource>
            ip-dscp-port-if-match-enable 5
```

9.3.5 Associating a DSCP or dot1p classification policy

A DSCP or dot1p classification policy must be associated with a SAP ingress QoS policy or an Ethernet port to map flows to an FC and profile for IP traffic received on SAP ingress.

The **dscp-classification** *policy-id* or **dot1p-classification** *policy-id* value identifies which classification policy is used to match IP packets and map the IP DSCP or dot1p to an FC and profile. The **no** form of the **dscp-classification** or **dot1p-classification** command associates the default classification policy (policy 1). The default policy maps all IP DSCP or dot1p values to FC "be" and profile "out". If a packet does not match any explicitly configured criteria in the policy, the **default-fc** mapping is used. For details about Layer 2 and Layer 3 scenarios, see [IP DSCP and dot1p classification policy support](#).

9.3.5.1 Associating a classification policy with a SAP ingress QoS policy

The **num-qos-classifiers** command allocates meters from the resources allocated toward the **qos-sap-ingress** pool, from the **ingress-internal-tcam** resource pool.

MAC, IPv4, and IPv6 criteria do not need to be defined because DSCP classification or dot1p classification is being used.

Use the following syntax to associate an IP DSCP classification policy with a SAP ingress policy.

```
configure> qos> sap-ingress policy-id [create]
      num-qos-classifiers num-resources [ipv6 | no-ipv6]
      [no] dscp-classification policy-id
      default-fc fc
```

The following example shows the command usage to associate an IP DSCP classification policy with a SAP ingress policy.

Example:

```
configure> qos> sap-ingress 1000 create
      num-qos-classifiers 16
      dscp-classification 101
      default-fc be
      exit
```

Use the following syntax to associate a dot1p classification policy with a SAP ingress policy.

```
configure> qos> sap-ingress policy-id [create]
      num-qos-classifiers num-resources [ipv6 | no-ipv6]
      [no] dot1p-classification policy-id
      default-fc fc
```

Example:

The following example shows the command usage to associate a dot1p classification policy with a SAP ingress policy.

```
- configure> qos> sap-ingress 1000 create
-   num-qos-classifiers 16
-   dscp-classification 101
-   default-fc be
-   exit
```

Use the following syntax to configure the classification policy to classify traffic to a forwarding class.

```
configure> qos> sap-ingress table-classification-criteria table-classification-criteria
```

Example:

The following example shows the command usage to configure the classification policy to classify traffic to a forwarding class.

```
configure> qos> sap-ingress table-classification-criteria none
```

9.3.5.2 Associating a classification policy with an Ethernet port

Users associate an IP DSCP classification policy with an Ethernet port using the **configure>port>ethernet>access>ingress>dscp-classification** command.

Users associate a dot1p classification policy with an Ethernet port using the **configure>port>ethernet>access>ingress>dot1p-classification** command.

Use the following syntax to associate an IP DSCP classification policy with an Ethernet port.

```
configure> port port-id
  ethernet> access> ingress
    [no] dscp-classification policy-id>
    [no] untagged-fc fc
```

Example

The following example shows the command usage to associate an IP DSCP classification policy with an Ethernet port.

Example:

```
configure> port 1/1/1
  ethernet
    access
      ingress
        dscp-classification 101
        untagged-fc 12
      exit
    exit
  exit
```

Use the following syntax to associate a dot1p classification policy with an Ethernet port.

```
- configure> port port-id
- ethernet> access> ingress
- [no] dot1p-classification policy-id>
- [no] untagged-fc fc
```

Example

Example:

```
configure> port 1/1/1
  ethernet
    access
      ingress
        dot1p-classification 101
        untagged-fc
      exit
    exit
  exit
```

9.3.6 Assigning and enabling DSCP and dot1p classification policies to a SAP

The 7210 SAS-Mxp supports table-based classification to assign an initial FC and profile on SAP ingress for epipe and VPLS SAPs, VPRN and IES interface SAPs, and RVPLS SAPs. The use of table-based classification (IP DSCP or dot1p) and CAM-based classification are mutually exclusive. That is, when table-based classification is used, any CAM-based classification configured from the SAP ingress QoS policy is ignored.

9.3.6.1 Assigning and enabling policies to Epipe and VPLS SAPs

Within Epipe and VPLS services, SAPs can be configured with an IP DSCP or dot1p classification policy per SAP. This applies to SAPs configured on an access port and on a hybrid port. Using the **enable-table-classification** command means the SAP uses table-based policies along with the meters defined in the SAP ingress policy. CAM-based resources from the SAP ingress policy are ignored.

The **enable-table-classification** command enables the use of IP DSCP or dot1p tables per SAP ingress to assign an FC and profile. Using table-based classification means ignoring CAM classification in the service ingress policy, using only meters from service ingress policy, and using the IP DSCP classification policy or dot1p classification policy that is configured in the SAP ingress policy. The default FC is assigned per SAP.

The **num-qos-classifiers** command allocates meters from the IFP, with resources taken from the ingress-internal-tcam resource pool toward qos-sap-ingress.

The **dscp-classification** command configures which classification policy is used to match IP packets and to map an IP DSCP value to an FC and profile.

The **dot1p-classification** command configures which classification policy is used to match IP packets and to map a dot1p value to an FC and profile.

The **table-classification-criteria** command provides an option for all traffic to use either dot1p classification, or DSCP classification, or both IP DSCP and Dot1p classification, or assign **default-fc** (**none** option) to all traffic. The **default-fc** command configures the FC assigned to all untagged packets. All untagged packets are mapped to the profile "out".

MAC, IPv4, and IPv6 criteria do not need to be defined because DSCP or dot1p classification is being used.

The CLI syntax below shows how to enable table classification for an Epipe and a VPLS service.

```
configure> service> epipe> sap> ingress>  
configure> service> vpls> sap> ingress>  
[no] qos policy-id [enable-table-classification]
```

The following is a sample SAP ingress QoS policy 1000 configured with DSCP classification policy 101 and default-fc "be". It is followed by a sample Epipe ingress SAP configured to use SAP ingress policy 1000 with table-based classification enabled.

Example:

```
configure> qos> sap-ingress 1000 create  
    num-qos-classifiers 16  
    fc af  
        meter 1  
        unknown-meter 2  
        multicast-meter 3  
        broadcast-meter 4  
    exit  
    dscp-classification 101  
    table-classification-criteria both-dscp-dot1p  
    default-fc be  
exit
```

Example:

```
configure> service> epipe> sap>
```

```
ingress
  qos 1000 [enable-table-classification]
exit
```

9.3.6.2 Assigning and enabling policies to IES and VPRN interface SAPs

Within IES and VPRN services, SAPs can be configured with an IP DSCP classification policy per SAP. This applies to SAPs configured on an access port and on a hybrid port. Using the **enable-table-classification** command means the SAP uses table-based policies along with the meters defined in the SAP ingress policy.

The **enable-table-classification**, **num-qos-classifiers**, **dscp-classification**, and **default-fc** commands for IES and VPRN interface SAPs operate similarly to Epipe and VPLS SAPs (see [Assigning and enabling policies to Epipe and VPLS SAPs](#)).

MAC, IPv4, and IPv6 criteria do not need to be configured because DSCP classification is being used.

Use the following syntax to enable table-based classification for an IES and a VPRN service.

```
configure>service>ies>interface>sap>ingress>
configure>service>vprn>interface>sap>ingress>
[no] qos policy-id [enable-table-classification]
```

The following is a sample SAP ingress QoS policy 1000 configured with DSCP classification policy 101 and **default-fc** "be". It is followed by a sample ingress SAP on an IES service interface configured to enable table-based classification using SAP ingress policy 1000.

Example:

```
configure> qos> sap-ingress 1000 create
  num-qos-classifiers 16
  fc af
    meter 1
    unknown-meter 2
    multicast-meter 3
    broadcast-meter 4
  exit
  dscp-classification 101
  default-fc be
  exit
```

Example:

```
configure>service>ies>interface>sap>
configure>service>vprn>interface>sap>
  ingress
    qos 1000 [enable-table-classification]
  exit
```

9.3.6.3 Assigning policies to RVPLS SAPs

For RVPLS SAPs configured on an access port, the 7210 SAS-Mxp supports RVPLS service with per-port IP DSCP classification policies or dot1p classification policies for bridged traffic received on SAPs configured in the RVPLS service. For routed traffic, per-IP interface IP DSCP or dot1p classification policies (that is, the QoS override policy) are used.

For RVPLS SAPs configured on a hybrid port, the network QoS policy of type "port" associated with network port ingress is used for RVPLS SAP bridged traffic classification and profile. Only the traffic classification will be used from the network policy. Meters are still used from the SAP ingress policy attached to the RVPLS SAP. For routed traffic received on a hybrid port, the IP DSCP or dot1p policy (that is, the QoS override policy) associated with the RVPLS IP interface is used for traffic classification and profile.

Only meters configured in the SAP ingress policy associated with RVPLS SAPs are used when table-based classification is enabled under the SAP associated with an RVPLS service.



Note:

All SAPs in an RVPLS service can use either table-based classification or CAM-based entries. There cannot be a mix of SAPs, where some SAPs use table-based classification and others use CAM entries.

9.3.6.3.1 Create a SAP ingress policy and assign the policy to an RVPLS SAP

The following examples create and assign a SAP ingress QoS policy to an RVPLS SAP. Table-based classification is enabled in the override policy associated with the IES interface that is associated with the RVPLS service. In this case, only meters from the SAP ingress QoS policy are used. Ingress CAM entries are ignored (not used).

Output example:

```
configure> qos> sap-ingress 1000 create
      num-qos-classifiers 16
      fc af
        meter 1
        unknown-meter 2
        multicast-meter 3
        broadcast-meter 4
      exit
exit
```

Output example:

```
configure> service> vpls (type rvpls)>
      service-name 'rvpls-example-svc'
      sap 1/1/1:100 create
        ingress
          qos 1000 [enable-table-classification]
        exit
      no shutdown
exit
```

9.3.6.3.2 Table-based classification per IP Interface for routed packets

For routed packets, although the IP DSCP classification is based on the DSCP policy that is attached to the IP interface, the **enable-table-classification** command must also be set on RVPLS SAPs for table-based classification to work correctly. If **enable-table-classification** is not set on an RVPLS SAP, only the profile will be taken from the routed-override-qos policy for that SAP. In this case, traffic classification (in accordance with TCAM-based classification) and meters will be taken from the SAP ingress policy attached to the RVPLS SAP.

The following syntax enables table-based classification and specifies the QoS override classification policy in the IES or VPRN interface RVPLS configuration. The *policy-id* specified in the **routed-override-qos-policy** command identifies the DSCP policy configured using the **configure>qos>dscp-classification** command.

```
configure> service> ies> interface> vpls service-name
    ingress
        [no] enable-table-classification
        [no] routed-override-qos-policy policy-id
    exit
exit
```

Example:

```
configure> service> ies> interface>
    vpls "rvpls-example-svc"
    ingress
        enable-table-classification
        routed-override-qos-policy 101
    exit
exit
exit
```

9.3.6.3.3 Table-based classification per port for bridged packets

For bridged packets, although the DSCP classification is based on the DSCP policy attached to the port, the **enable-table-classification** command must also be set in the IES or VPRN interface context as well as the respective RVPLS SAP context for table-based classification to work correctly (as shown in the Example). If **enable-table-classification** is not set on the respective RVPLS SAP then only profile will be taken from the port policy for that SAP. In this case, classification (in accordance with TCAM-based classification) and meters will be taken from the SAP ingress policy.

The following syntax enables table-based classification on an Ethernet port and specifies the DSCP classification policy in the **port>ethernet>access>ingress** command.

```
configure>port>
    ethernet>
        [no] enable-table-classification
    exit

configure>port>ethernet>access>ingress>
    [no] dscp-classification policy-id
    [no] untagged-fc fc
exit
```

Example

Example:

```
configure>port 1/1/1
    ethernet
        enable-table-classification
    exit

configure>port> 1/1/1
    ethernet>access>ingress>
        dscp-classification 101
```

```
        untagged-fc ef
    exit
exit

configure> service> ies> interface>
    vpls "rvpls-example-svc"
    ingress
        enable-table-classification
    exit
exit
exit
```

9.4 Service meter for SAP ingress (7210 SAS-Mxp)

Service meters for SAP ingress provide an option to use meter resources from the ingress service-meter pool, which provides a larger number of meters/policers for use by access SAPs. This option is available only with table-based classification; it is not available when CAM-based classification is used.



Note:

Table-based classification uses meters from either the TCAM pool or the service meter pool, based on the SAP ingress policy type. If the SAP ingress policy is configured to use the **use-svc-meter-pool** parameter, the policy uses the service meter pool, otherwise the policy uses the TCAM meter pool.

Access to larger numbers of meters/policers from the service-meter pool is useful when there is a need to enforce bandwidth limits for all the FCs and traffic-types (that is, unicast, broadcast, multicast, and unknown-unicast) across a large number of access SAPs. The following functionality is available with the service-meter option:

- All access SAPs that use the service-meter pool share a single FC-to-meter-map policy that defines the meter to use for the configured FCs (there is a single configurable policy other than the default FC-to-meter map).
- The SAP ingress policy must be configured to use table-based classification with both IP DSCP and dot1p classification policies configured.
- SAP ingress policies can be configured with different IP DSCP and dot1p classification policies, which provide flexibility when configuring mappings for IP DSCP and dot1p values to FC for different SAPs.
- Even though all access SAPs share the same fc-to-meter-map policy, the meter instances and counter instances are different for each SAP that is configured. This means that each SAP gets an individual instance of meter to enforce rate-limits for traffic that is received on the SAP.

The following usage restrictions apply to meters across FCs and traffic types:

- If a meter is mapped to one or more traffic types (unicast, broadcast, multicast, and unknown-unicast) in a specific FC, the meter must be used with the same mapping in other FCs.
- Resources are always allocated in chunks of 1, 2, 4, 8, 16, or 32 meters. Counters are allocated in chunks of power of 2.



Note:

Service meter is supported for access SAPs that are configured in Epipe, VPLS, IES, and VPRN services. Service meter is not supported for access SAPs that are configured in an RVPLS service.

9.4.1 Default service meter policy

Output example

The following is a sample default service meter policy output.

```
A:Dut-A>config>qos# info
-----
#-----
echo "QoS Policy Configuration"
#-----
      dscp-classification 1 create
      description "Default DSCP Classification policy"
      exit
-----
A:Dut-A>config>qos# fc-meter-map 1
A:Dut-A>config>qos>fc-meter-map# info
-----
-----
A:Dut-A>config>qos>fc-meter-map# info detail
-----
      counter-mode in-out-profile-count
      meter 1 create
      mode trtcm1
      adaptation-rule cir closest pir closest
      rate cir 0 pir max
      mbs default kbits
      cbs default kbits
      color-mode color-blind
      exit
      fc be create
      meter 1
      no broadcast-meter
      no multicast-meter
      no unknown-meter
      exit
-----
```

9.4.2 Resource usage for service meters

The number of meters used for a SAP ingress policy that is configured to use the service-meter pool is equal to the number of unique meters mapped to different FCs in the fc-meter-map policy that is associated with the SAP ingress policy.

The fc-meter-map defines the association between the meters and the FC-to-meter. Meter resources are allocated only to meters that have an associated FC. For example, if 10 meters are created in the fc-meter-map, and only 5 meters are associated with an FC, the system allocates only 5 meters per SAP (and not 10 meters) and rounds off the number of meters to the nearest power of 2, which results in eight (8) meters to be allocated. Meters that are not associated with an FC are ignored for resource allocation. The number of counters allocated is equal to twice the number of meters per SAP.



Note:

- Based on the number of meters that the SAP requires, the amount of meter resources allocated in hardware is equal to the nearest power of 2 greater than the number of meters required. For example, if the number of required meter resources as determined by the fc-meter-map is five (5), then eight (8) meter resources are allocated in hardware. Similarly, if the

number of meter resources required as determined by the `fc-meter-map` is 10, then 16 meter resources are allocated in hardware.

The number of counters allocated corresponds to the number of meters. If the number of counter resources required as determined by the `fc-meter-map` is 16, then 32 counter resources are allocated in the hardware.

- Even if no CAM resources are used for either classification (that is, matching of packet header QoS bits) or for meters, at least 1 slice must be allocated to SAP ingress classification and policing from the **ingress-internal-tcam>qos-sap-ingress-resource** pool.

The slice is used to allocate resources that are required for the default SAP ingress policy (which requires 1 tcam meter resource and 2 tcam classification resources), when the SAP is configured using the **create** command. Otherwise, SAP configuration using the **create** command will fail. The remaining SAP ingress resource pool CAM slices from the `ingress-internal-tcam` resource pool can be reassigned to other entities.

9.4.2.1 Examples for service meters with computation of resource usage

The examples in this section use the following terminology:

- `fc-name` (Ucast) - unicast meter that is used for FC and specified by the *fc-name*

For example, to configure a unicast meter that uses the FC value "be," the notation `be(Ucast)` is used.

- B - broadcast meter
- U - unknown-unicast meter (also referred to as unknown meter)
- M - multipoint meter

This section provides examples of usage for service meters.

Example:

Single FC configured to use the default unicast meter for all traffic (both unicast and BUM traffic)

```
*A:Dut-A>config>qos>fc-meter-map# info
-----
meter 1 create
exit
meter 2 create
exit
meter 3 create
exit
meter 4 create
exit
meter 11 create
exit
fc be create
exit
-----
Meters being used = 1 [Meter 1 for be(Ucast) and be(BUM)]
Counters being used = 2 (2 x num of meters)
Meter reserved per SAP = 1(nearest exponent of 2 equal to 1)
Counters reserved per SAP = 2(2 x num of meters reserved per SAP)
=====
```

Example:

Single FC configured to use the default unicast meter for all unicast traffic and default multipoint meter (meter 11) for all BUM traffic

```
*A:Dut-A>config>qos>fc-meter-map# info
-----
meter 1 create
exit
meter 2 create
exit
meter 3 create
exit
meter 4 create
exit
meter 11 multipoint create
exit
fc be create
exit
-----
Meters being used = 2 [Meter 1 for be(Ucast) + Meter 11 for be(BUM)]
Counters being used = 4 (2 x num of meters)
Meter reserved per SAP = 2 (nearest exponent of 2 equal to 2)
Counters reserved per SAP = 4(2 x num of meters reserved per SAP)
=====
```

Example:

Two FCs configured to use the default unicast meter for all unicast traffic and default multipoint meter (meter11) for all BUM traffic

```
*A:Dut-A>config>qos>fc-meter-map# info
-----
meter 1 create
exit
meter 2 create
exit
meter 3 create
exit
meter 4 create
exit
meter 11 multipoint create
exit
fc be create
exit
fc l2 create
exit
-----
Meters being used = 2 [Meter 1 for be (Ucast) and l2 (Ucast) + Meter 11 for be(BUM)
and l2(BUM)]
Counters being used = 4 (2 x num of meters)
Meter reserved per SAP = 2(nearest exponent of 2 equal to 2)
Counters reserved per SAP = 4(2 x num of meters reserved per SAP)
=====
```

Example:

Two FCs configured to use the default unicast meter for all unicast traffic, with FC "be" configured to use default multipoint meter (meter 11) for all BUM traffic, with FC "l2" configured to use an explicit

broadcast meter only broadcast traffic and the default multipoint meter 11 for both multicast and unknown-unicast traffic

```
*A:Dut-A>config>qos>fc-meter-map# info
-----
meter 1 create
exit
meter 2 create
exit
meter 3 create
exit
meter 4 create
exit
meter 11 multipoint create
exit
fc be create
exit
fc l2 create
    broadcast-meter 2
exit
-----
Meters being used = 3 [Meter 1 for be(Ucast) and l2( Ucast) + Meter 11 for be(BUM)
and l2 (UM) + Meter 2 for l2(B)]
Counters being used = 6 (2 x num of meters)
Meter reserved per SAP = 4(nearest exponent of 2 greater than 3)
Counters reserved per SAP = 8(2 x num of meters reserved per SAP)
=====
```

Example:

Two FCs configured, with one FC (that is, FC "be") to use the explicitly configured unicast meter for unicast traffic and explicit meter for BUM traffic, and the other FC (that is, FC "l2") configured to use default unicast meter for unicast traffic and explicit meter for BUM traffic

```
*A:Dut-A>config>qos>fc-meter-map# info
-----
meter 1 create
exit
meter 2 create
exit
meter 3 create
exit
meter 4 create
exit
meter 11 multipoint create
exit
fc be create
    meter 2
    broadcast-meter 2
    multicast-meter 2
    unknown-meter 2
exit
fc l2 create
    broadcast-meter 2
    multicast-meter 2
    unknown-meter 2
exit
-----
Meters being used = 2 [Meter 1 for l2( Ucast) + Meter 2 for be(Ucast) , be(BUM)
and l2 (BUM)]
Counters being used = 4 (2 x num of meters)
Meter reserved per SAP = 2(nearest exponent of 2 equal to 2)
```

```
Counters reserved per SAP = 4(2 x num of meters reserved per SAP)
=====
```

Example:

One FC (that is, FC "be") configured to use the explicitly configured unicast meter for all unicast traffic and BUM traffic

```
*A:Dut-A>config>qos>fc-meter-map# info
-----
meter 1 create
exit
meter 2 create
exit
meter 3 create
exit
meter 4 create
exit
fc be create
    meter 2
exit
-----
Meters being used = 1 [Meter 2 for be(Ucast) and be(BUM)]
Counters being used = 2 (2 x num of meters)
Meter reserved per SAP = 1(nearest exponent of 2 equal to 1)
Counters reserved per SAP = 2(2 x num of meters reserved per SAP)
=====
```

9.5 Calculating resources required for classification

This section provides examples for calculating the resources required for SAP-ingress policy classification when using CAM-based classification and table-based classification.

9.5.1 Examples: calculating resources required for CAM-based classification

This section provides examples for calculating the amount of resources needed for a service ingress policy when CAM-based classification is used. For calculations when IP DSCP table-based classification is used, see [Examples: calculating resources required for IP DSCP table-based classification with CAM-based policing \(7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12\)](#).

The resource calculation shown for VLL is also applicable for VPRN services.

The examples in this section use the two equations below to calculate the value for **num-qos-classifiers** used in the sap-ingress QoS policy. See [Computation of resources used per SAP ingress policy for CAM-based classification](#) for more information about these equations.

- total number of classification entries (TC)

$$TC = \# E(i) * C(i)$$

i = nc, h1, ef, h2, l1, af, l2, be

where:

- E(i) is the number of match-criteria entries that classify packets to FC_i

- $C(i)$ is the number of classification entries that are required by FC_i to identify different traffic types
- total number of QoS resources required (TQ)

$$TQ = \max [(TC), (2 * TP)]$$

where:

- TP is the number of meters/policers used

In addition, the examples show how to determine the number of classification entries for each forwarding class. For example, FCh2 (shown below) is the sum of four traffic types: unicast (U), broadcast (B), multicast (M), and unknown-unicast (U-u). See [Calculating the number of classification entries per FC](#) for more information.

$$FCh2 = 1 + 0 + 1 + 0 = 2 \text{ (generalized to } FCh2 = U + B + M + U-u \text{)}$$

If BUM entries are not explicit and multipoint traffic is expected, meter "11" is used and the "M" traffic type is given a "1".

9.5.1.1 Example 1

Example

```
sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 8
  meter 1 create
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  scope template
  default-fc be
  fc be create
    meter 3
  exit
  fc af create
    meter 1
  exit
  fc l1 create
    meter 3
  exit
  fc h2 create
    meter 3
  exit
  mac-criteria dot1p-only
  entry 1 create
    match dot1p 7
    action fc af
  exit
  entry 2 create
    match dot1p 5
    action fc l1
```

```

        exit
    entry 3 create
        match dot1p 6
        action fc h2
    exit
exit
exit

```

In the preceding example, assuming the policy is attached to a SAP in a VPLS service, compute the number of classification entries per FC as follows.

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 0 + 0 + 0 + 0 = 0 \quad FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 0 + 1 + 0 = 2$$

Because FCh2 uses unicast meter, an entry is needed to identify unicast traffic type explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter "11".

$$FCI1 = 1 + 0 + 1 + 0 = 2 \quad FCaf = 1 + 0 + 1 + 0 = 2 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 1 + 0 = 2$$

Using the equation, calculate the total number of classification entries (TC) used by this policy, as follows:

$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (1 * 2)h2 + (1 * 2)l1 + (1 * 2)af + (0 * 0)l2 + (1 * 2)be = 8$ (because three explicit match criteria entries are used to map traffic to each of FC H2, FC L1, and FC AF along with a default classification entry for FC BE).

The total number of meters used = 3 (because FCs use meter "1", meter "3" and meter "11").

In this example, **num-qos-classifiers 8** is used (maximum of $(8, (2 * 3))$).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following computation is made:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 0 + 0 + 0 + 0 = 0 \quad FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1 \quad FCaf = 1 + 0 + 0 + 0 = 1 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation, total classification entries used = 4 and meters used = 2.

As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Hence, it is recommended to create a different policy with the required number of resources (that is, with **num-qos-classifiers 4**)

9.5.1.2 Example 1a (default multipoint meter 11 is not used):

Example

```

sap-ingress 10 create
    description "example-policy"
    num-qos-classifiers 4
    meter 1 create
        rate cir 0 pir max
    exit
    meter 3 create
        rate cir 100 pir 100
    exit
scope template

```

```

default-fc be

fc be create
    meter 3
exit
fc af create
    meter 1
exit
fc l1 create
    meter 3
exit
fc h2 create
    meter 3
exit
mac-criteria dot1p-only
entry 1 create
    match dot1p 7
    action fc af
exit
entry 2 create
    match dot1p 5
    action fc l1
exit
entry 3 create
    match dot1p 6
    action fc h2
exit
exit
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, compute the number of classification entries per FC as follows:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 0 + 0 + 0 + 0 = 0 \quad FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 0 + 0 + 0 = 1$$

Because FCh2 uses unicast meter for all traffic types, we need an entry to classify all traffic types to FCh2 explicitly.

$$FCI1 = 1 + 0 + 0 + 0 = 1 \quad FCaf = 1 + 0 + 0 + 0 = 1 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 0 + 0 = 1$$

Using the equation, calculate the total classification entries used by this policy, as follows:

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (1 * 1)h2 + (1 * 1)l1 + (1 * 1)af + (0 * 0)l2 + (1 * 1)be = 4$$

(because three explicit match criteria entries are used to map traffic to each of FC H2, FC L1, and FC AF along with a default classification entry for FC BE).

The total number of meters used = 2 (because FCs use meter "1" and meter "3").

In this example, **num-qos-classifiers 4** is used (maximum of (4, (2 * 2))). Use of unicast meter for all traffic-types allows for use QoS resources efficiently.

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 0 + 0 + 0 + 0 = 0 \quad FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 0 + 0 + 0 = 1 \\ FCI1 = 1 + 0 + 0 + 0 = 1 \quad FCaf = 1 + 0 + 0 + 0 = 1 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation for TC calculation, total classification entries used = 4 and meters used = 2.

As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Hence, it is recommended to create a different policy with the required number of resources (that is, with **num-qos-classifiers 4**).

9.5.1.3 Example 2

Example

```
sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 16
  meter 1 create
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 2 create
    rate cir 1 pir 20
  exit
  scope template
  default-fc be
  fc be create
    meter 3
    broadcast-meter 2
  exit
  fc af create
    meter 3
    broadcast-meter 2
  exit
  fc ll create
    meter 3
    broadcast-meter 2
  exit
  fc h2 create
    meter 3
    broadcast-meter 2
  exit
  mac-criteria dot1p-only
    entry 1 create
      match dot1p 7
      action fc af
    exit
    entry 2 create
      match dot1p 5
      action fc ll
    exit
    entry 3 create
      match dot1p 6
      action fc h2
    exit
  exit
exit
```

In the preceding example, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC as:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 0 + 0 + 0 + 0 = 0 \quad FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 1 + 1 + 0 = 3$$

Because FCh2 uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Another entry is needed to classify multicast and unknown-unicast traffic type to FCh2 and use the default meter "11".

$$FCI1 = 1 + 1 + 1 + 0 = 3 \quad FCaf = 1 + 1 + 1 + 0 = 3 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 1 + 1 + 0 = 3$$

Using the above equation for TC calculation, to get the total classification entries used = 12 (because three explicit match criteria entries map to each of FC H2, L1, and AF along with a default classification rule for BE).

The number of meters used = 3 (because FCs use only meter "2", meter "3" and meter "11").

Hence, in this example **num-qos-classifiers 16** is used (that is, maximum of (12, (2*3))).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 0 + 0 + 0 + 0 = 0 \quad FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1 \quad FCaf = 1 + 0 + 0 + 0 = 1 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation, to get total classification entries used = 4 and Meters used = 1.

As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Hence, it is recommended to create a different policy with the required number of resources (that is, with **num-qos-classifiers 4**)

9.5.1.4 Example 2a (default multipoint meter "11" is not used)

Example

```
sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 8

  meter 1 create
    rate cir 0 pir max
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 2 create
    rate cir 1 pir 20
  exit
  scope template
  default-fc be
  fc be create
    meter 3
    broadcast-meter 2
  exit
  fc af create
    meter 3
```

```

        broadcast-meter 2
    exit
    fc l1 create
        meter 3
        broadcast-meter 2
    exit
    fc h2 create
        meter 3
        broadcast-meter 2
    exit
    mac-criteria dot1p-only
    entry 1 create
        match dot1p 7
        action fc af
    exit
    entry 2 create
        match dot1p 5
        action fc l1
    exit
    entry 3 create
        match dot1p 6
        action fc h2
    exit
exit

```

In the preceding example, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC as:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 0 + 0 + 0 + 0 = 0 \quad FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 0 + 1 + 0 = 2$$

Because FCh2 uses unicast meter for unicast, multicast, and unknown-unicast traffic, and broadcast meter for broadcast traffic, two entries are needed.

$$FCI1 = 1 + 0 + 1 + 0 = 2 \quad FCaf = 1 + 0 + 1 + 0 = 2 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 1 + 0 = 2$$

Using the above equation, to get the total classification entries used = 8 (since three explicit match criteria entries map to each of FC H2, L1, and AF along with a default classification rule for BE).

The number of meters used = 2 (because FCs use only meter "2" and meter "3").

Hence, in this example **num-qos-classifiers 8** is used (that is, maximum of (8, (2*2))).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 0 + 0 + 0 + 0 = 0 \quad FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1 \quad FCaf = 1 + 0 + 0 + 0 = 1 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation, to get total classification entries used = 4 and Meters used = 1.

As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Hence, it is recommended to create a different policy with the required number of resources (that is, with **num-qos-classifiers 4**)

9.5.1.5 Example 3

Example

```
sap-ingress 10 create
  description "example-policy-2"
  num-qos-classifiers 16
  meter 1 create
    rate cir 100 pir 100
  exit
  meter 11 multipoint create
    rate cir 1 pir 20
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 2 create
    rate cir 1 pir 20
  exit
  meter 4 create
    rate cir 10 pir 100
  exit
  meter 5 create
    rate cir 10 pir 10
  exit
  scope template
  default-fc be
  fc af create
    meter 3
    broadcast-meter 2
    multicast-meter 4
  exit
  fc l1 create
    meter 3
    broadcast-meter 2
  exit
  fc h2 create
    meter 3
    broadcast-meter 2
  exit
  fc h1 create
    meter 5
    broadcast-meter 4
    multicast-meter 4
    unknown-meter 4
  exit
  mac-criteria dot1p-only
    entry 1 create
      match dot1p 7
      action fc af
    exit
    entry 2 create
      match dot1p 5
      action fc l1
    exit
    entry 3 create
      match dot1p 6
      action fc h2
    exit
    entry 4 create
      match dot1p 3
```

```

        action fc h1
        exit
    exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the classification entries used per FC are:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 1 + 1 + 1 + 1 = 4$$

Because FCh1 uses unicast, broadcast, multicast and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 1 + 1 + 0 = 3$$

Because FCh2 uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Another entry is needed to classify multicast and unknown-unicast traffic type to the same FC and use the default meter "11".

$$FCI1 = 1 + 1 + 1 + 0 = 3$$

Because FCI1 uses only unicast meter, an entry is needed to identify this traffic type explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter "11".

$$FCaf = 1 + 1 + 1 + 0 = 3$$

Because FCaf uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 1 + 0 = 2$$

Using the above equation, the total classification entries used = 15 and meters used = 6.

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following results:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 1 + 0 + 0 + 0 = 1 \quad FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1 \quad FCaf = 1 + 0 + 0 + 0 = 1 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation, the total classification entries used = 5 and meters used = 3 (because all FCs used only meter "1", meter "3" and meter "5").

9.5.1.6 Example 3a (default multipoint meter "11" is not used):

Example

```

sap-ingress 10 create
description "example-policy-2"
num-qos-classifiers 12
meter 1 create
rate cir 100 pir 100

```

```

exit
meter 3 create
    rate cir 100 pir 100
exit
meter 2 create
    rate cir 1 pir 20
exit
meter 4 create
    rate cir 10 pir 100
exit
meter 5 create
    rate cir 10 pir 10
exit
scope template
default-fc be
fc af create
    meter 3
    broadcast-meter 2
    multicast-meter 4
exit
fc l1 create
    meter 3
    broadcast-meter 2
exit
fc h2 create
    meter 3
    broadcast-meter 2
exit
fc h1 create
    meter 5
    broadcast-meter 4
    multicast-meter 4
    unknown-meter 4
exit
mac-criteria dot1p-only
entry 1 create
    match dot1p 7
    action fc af
exit
entry 2 create
    match dot1p 5
    action fc l1
exit
entry 3 create
    match dot1p 6
    action fc h2
exit
entry 4 create
    match dot1p 3
    action fc h1
exit
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the classification entries used per FC are:

$$FCnc = 0 + 0 + 0 + 0 = 0 \quad FCh1 = 1 + 1 + 1 + 1 = 4$$

Because FCh1 uses unicast, broadcast, multicast and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$FCef = 0 + 0 + 0 + 0 = 0 \quad FCh2 = 1 + 0 + 1 + 0 = 2$$

Because FCh2 uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly, multicast and unknown-unicast traffic use the same resource as the unicast traffic.

$$FCI1 = 1 + 0 + 1 + 0 = 2$$

Because FCI1 uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. multicast and unknown-unicast traffic use the same resource as the unicast traffic.

$$FCaf = 1 + 1 + 1 + 0 = 3$$

Because FCaf uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 0 + 0 = 1$$

Because no explicit meters are configured for FC "be", it uses meter "1" for all traffic types and needs one entry is needed to identify these traffic types.

Using the above equation, the total classification entries used = 12 and meters used = 5. The **num-qos-classifiers** can be set to 12 (the minimum value).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following results:

$$\begin{aligned} FCnc &= 0 + 0 + 0 + 0 = 0 & FCh1 &= 1 + 0 + 0 + 0 = 1 & FCef &= 0 + 0 + 0 + 0 = 0 & FCh2 &= 1 + 0 + 0 + 0 = 1 \\ FCI1 &= 1 + 0 + 0 + 0 = 1 & FCaf &= 1 + 0 + 0 + 0 = 1 & FCI2 &= 0 + 0 + 0 + 0 = 0 & FCbe &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the above equation, the total classification entries used = 5 and meters used = 3 (because all FCs used only meter "1", meter "3" and meter "5"). For Epipe service a policy with **num-qos-classifiers** set to 6 can be used.

9.5.1.7 Example 4

Example

```
sap-ingress 10 create
  description "example-policy-3"
  num-qos-classifiers 32
  meter 1 create
    rate cir 100 pir 100
  exit
  meter 11 multipoint create
    rate cir 1 pir 20
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 2 create
    rate cir 1 pir 20
  exit
  meter 4 create
    rate cir 10 pir 100
  exit
```

```
meter 5 create
    rate cir 10 pir 10
exit
meter 6 create
    rate cir 11 pir 100
exit
meter 8 create
    rate cir 20 pir 100
exit
scope template
default-fc be
fc af create
    meter 3
    broadcast-meter 2
    multicast-meter 4
exit
fc l1 create
    meter 3
    broadcast-meter 2
exit
fc h2 create
    meter 3
    broadcast-meter 2
exit
fc h1 create
    meter 5
    broadcast-meter 4
    multicast-meter 4
    unknown-meter 4
exit
fc ef create
    meter 6
    broadcast-meter 2
    multicast-meter 8
exit
fc nc create
    meter 6
    broadcast-meter 2
    multicast-meter 8
exit
mac-criteria dot1p-only
entry 1 create
    match dot1p      4
    action fc af
exit
entry 2 create
    match dot1p      5
    action fc l1
exit
entry 3 create
    match dot1p      6
    action fc h2
exit
entry 4 create
    match dot1p      3
    action fc h1
exit
entry 5 create
    match dot1p      2
    action fc ef
exit
entry 6 create
    match dot1p      7
    action fc nc
```

```

        exit
      exit
    exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, compute the classification entries per FC as:

$$FCnc = 1 + 1 + 1 + 0 = 3$$

Because FCnc uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh1 = 1 + 1 + 1 + 1 = 4$$

Because FCh1 uses unicast, broadcast, multicast and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$FCef = 1 + 1 + 1 + 0 = 3$$

Because FCef uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh2 = 1 + 1 + 1 + 0 = 3$$

Because FCh2 uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Another entry is needed to classify multicast and unknown-unicast traffic type to the same FC and use the default meter "11".

$$FCI1 = 1 + 1 + 1 + 0 = 3 \quad FCaf = 1 + 1 + 1 + 0 = 3$$

Because FCaf uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 1 + 0 = 2$$

Using the above equation, the total classification entries used = 21 and meters used = 8.

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$FCnc = 1 + 0 + 0 + 0 = 1 \quad FCh1 = 1 + 0 + 0 + 0 = 1 \quad FCef = 1 + 0 + 0 + 0 = 1 \quad FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1 \quad FCaf = 1 + 0 + 0 + 0 = 1 \quad FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation, the total classification entries used = 7 and meters used = 4.

As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Hence, it is recommended to create a different policy with the required number of resources (that is, with **num-qos-classifiers 8**)

9.5.1.8 Example 4a (default multipoint meter "11" is not used):

Example

```
sap-ingress 10 create
  description "example-policy-3"
  num-qos-classifiers 20
  meter 1 create
    rate cir 100 pir 100
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 2 create
    rate cir 1 pir 20
  exit
  meter 4 create
    rate cir 10 pir 100
  exit
  meter 5 create
    rate cir 10 pir 10
  exit
  meter 6 create
    rate cir 11 pir 100
  exit
  meter 8 create
    rate cir 20 pir 100
  exit

scope template

default-fc be
fc af create
  meter 3
  broadcast-meter 2
  multicast-meter 4
exit
fc l1 create
  meter 3
  broadcast-meter 2
exit
fc h2 create
  meter 3
  broadcast-meter 2
exit
fc h1 create
  meter 5
  broadcast-meter 4
  multicast-meter 4
  unknown-meter 4
exit
fc ef create
  meter 6
  broadcast-meter 2
  multicast-meter 8
exit
fc nc create
  meter 6
  broadcast-meter 2
  multicast-meter 8
exit
```

```

mac-criteria dot1p-only
entry 1 create
    match dot1p 4
    action fc af
exit
entry 2 create
    match dot1p 5
    action fc l1
exit
entry 3 create
    match dot1p 6
    action fc h2
exit
entry 4 create
    match dot1p 3
    action fc h1
exit
entry 5 create
    match dot1p 2
    action fc ef
exit
entry 6 create
    match dot1p 7
    action fc nc
exit
exit
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, compute the classification entries per FC as:

$$FCnc = 1 + 1 + 1 + 0 = 3$$

Because FCnc uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh1 = 1 + 1 + 1 + 1 = 4$$

Because FCh1 uses unicast, broadcast, multicast and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$FCef = 1 + 1 + 1 + 0 = 3$$

Because FCef uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh2 = 1 + 1 + 1 + 0 = 3$$

Because FCh2 uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. multicast and unknown-unicast traffic of the same FC use the unicast resources (both meter and classification entry).

$$FCI1 = 1 + 1 + 1 + 0 = 3 \quad FCaf = 1 + 1 + 1 + 0 = 3$$

Because FCaf uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCI2 = 0 + 0 + 0 + 0 = 0 \quad FCbe = 1 + 0 + 0 + 0 = 1$$

Because FCbe uses a single meter for all traffic-types only a single meter and single entry is needed.

Using the above equation, the total classification entries used = 20 and meters used = 7, **num-qos-classifiers** to use is 20 (the minimum value).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$\begin{aligned} FCnc &= 1 + 0 + 0 + 0 = 1 & FCh1 &= 1 + 0 + 0 + 0 = 1 & FCef &= 1 + 0 + 0 + 0 = 1 & FCh2 &= 1 + 0 + 0 + 0 = 1 \\ FCI1 &= 1 + 0 + 0 + 0 = 1 & FCaf &= 1 + 0 + 0 + 0 = 1 & FCI2 &= 0 + 0 + 0 + 0 = 0 & FCbe &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the above equation, the total classification entries used = 7 and meters used = 4.

As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Hence, it is recommended to create a different policy with the required number of resources (that is, with **num-qos-classifiers 8**).

9.5.1.9 Example 5

Example

```
sap-ingress 10 create
  description "example-policy-3"
  num-qos-classifiers 32
  meter 1 create
    rate cir 100 pir 100
  exit
  meter 11 multipoint create
    rate cir 1 pir 20
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 2 create
    rate cir 1 pir 20
  exit
  meter 4 create
    rate cir 10 pir 100
  exit
  meter 5 create
    rate cir 10 pir 10
  exit
  meter 6 create
    rate cir 11 pir 100
  exit
  meter 8 create
    rate cir 20 pir 100
  exit
  scope template
  default-fc be
  fc af create
```

```

        meter 3
        broadcast-meter 2
        multicast-meter 4
    exit
    fc l1 create
        meter 3
        broadcast-meter 2
    exit
    fc h2 create
        meter 3
        broadcast-meter 2
    exit
    fc h1 create
        meter 5
        broadcast-meter 4
        multicast-meter 4
        unknown-meter 4
    exit
    fc ef create
    exit
    fc nc create
        meter 6
        broadcast-meter 2
        multicast-meter 8
    exit
    mac-criteria dot1p-only
        entry 1 create
            match dot1p 4
            action fc af
        exit
        entry 2 create
            match dot1p 5
            action fc l1
        exit
        entry 3 create
            match dot1p 6
            action fc h2
        exit
        entry 4 create
            match dot1p 3
            action fc h1
        exit
        entry 5 create
            match dot1p 2
            action fc ef
        exit
        entry 6 create
            match dot1p 7
            action fc nc
        exit
    exit
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, get the classification entries used per FC:

$$FCnc = 1 + 1 + 1 + 0 = 3$$

Because FCnc uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh1 = 1 + 1 + 1 + 1 = 4$$

Because FCh1 uses unicast, broadcast, multicast and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$FCef = 1 + 0 + 1 + 0 = 2$$

Because no meters are explicitly configured, FCef uses the appropriate default meters all the traffic types (that is, unicast traffic uses unicast meter "1" and broadcast, multicast, and unknown-unicast traffic uses multipoint meter "11").

$$FCh2 = 1 + 1 + 1 + 0 = 3$$

Because FCh2 uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Another entry is needed to classify multicast and unknown-unicast traffic type to the same FC and use the default meter "11".

$$FCI1 = 1 + 1 + 1 + 0 = 3 \quad FCaf = 1 + 1 + 1 + 0 = 3$$

Because FCaf uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$\begin{aligned} FCl2 &= 0 + 0 + 0 + 0 = 0 \\ FCbe &= 1 + 0 + 1 + 0 = 2 \end{aligned}$$

Using the above equation, the total classification entries used = 20 and meters used = 8.

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$\begin{aligned} FCnc &= 1 + 0 + 0 + 0 = 1 & FCh1 &= 1 + 0 + 0 + 0 = 1 & FCef &= 1 + 0 + 0 + 0 = 1 & FCh2 &= 1 + 0 + 0 + 0 = 1 \\ FCI1 &= 1 + 0 + 0 + 0 = 1 & FCaf &= 1 + 0 + 0 + 0 = 1 & FCl2 &= 0 + 0 + 0 + 0 = 0 & FCbe &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the above equation, to get the total classification entries used = 7 and meters used = 4.

9.5.1.10 Example 6

Example

```
sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 16

  meter 1 create
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
```

```
meter 4 create
    rate cir 10 pir 50
exit

scope template

default-fc be
fc be create
    meter 3
exit
fc af create
    meter 1
exit
fc l1 create
    meter 3
    multicast-meter 4
exit
fc h2 create
    meter 3
exit

mac-criteria dot1p-only
    entry 1 create
        match dot1p 7
        action fc af
    exit
    entry 2 create
        match dot1p 5
        action fc l1
    exit
    entry 3 create
        match dot1p 6
        action fc h2
    exit
exit
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the following number of classification entries per FC:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 & \text{FCh1} &= 0 + 0 + 0 + 0 = 0 & \text{FCef} &= 0 + 0 + 0 + 0 = 0 & \text{FCh2} &= 1 + 0 + 1 + 0 = 2 \\ \text{FCI1} &= 1 + 0 + 1 + 0 = 2 \end{aligned}$$

Because FCI1 uses unicast meter and multicast meter, an entry is needed to identify these traffic types explicitly. Broadcast and unknown-unicast traffic is also classified using the same entry as multicast and use the same meter.

$$\text{FCaf} = 1 + 0 + 1 + 0 = 2$$

Because FCaf uses unicast meter, an entry is needed to identify these traffic types explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter "11".

$$\text{FCI2} = 0 + 0 + 0 + 0 = 0 \quad \text{FCbe} = 1 + 0 + 1 + 0 = 2$$

Using the above equation, the total classification entries used = 8 and meters used = 4.

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 & \text{FCh1} &= 0 + 0 + 0 + 0 = 0 & \text{FCef} &= 0 + 0 + 0 + 0 = 0 & \text{FCh2} &= 1 + 0 + 0 + 0 = 1 \\ \text{FCI1} &= 1 + 0 + 0 + 0 = 1 & \text{FCaf} &= 1 + 0 + 0 + 0 = 1 & \text{FCI2} &= 0 + 0 + 0 + 0 = 0 & \text{FCbe} &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the above equation, the total classification entries used = 4 and meters used = 2.

9.5.1.11 Example 7

Example

```
sap-ingress 10 create
  num-qos-classifiers 8
  meter 1 create
  exit
  meter 11 multipoint create
  exit
  meter 3 create
  exit
  meter 4 create
  exit
  fc be create
    meter 1
    broadcast-meter 11
    mulitcast-meter 4
  exit
  fc af create
    meter 3
  exit
  default-fc be
  match entry 1
    dot1p 7 fc af
  exit
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the following number of classification entries per FC are:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 & \text{FCh1} &= 0 + 0 + 0 + 0 = 0 & \text{FCef} &= 0 + 0 + 0 + 0 = 0 & \text{FCh2} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCI1} &= 0 + 0 + 0 + 0 = 0 & \text{FCaf} &= 1 + 0 + 1 + 0 = 2 \end{aligned}$$

Because FCaf uses unicast meter, an entry is needed to identify these traffic types explicitly. Another entry is needed entry to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter "11".

$$\text{FCI2} = 0 + 0 + 0 + 0 = 0 \quad \text{FCbe} = 1 + 1 + 1 + 0 = 3$$

Because FCbe uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

Using the above equation, the total classification entries used = 5 and meters used = 4.

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 & \text{FCh1} &= 0 + 0 + 0 + 0 = 0 & \text{FCef} &= 0 + 0 + 0 + 0 = 0 & \text{FCh2} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCI1} &= 0 + 0 + 0 + 0 = 0 & \text{FCaf} &= 1 + 0 + 0 + 0 = 1 & \text{FCI2} &= 0 + 0 + 0 + 0 = 0 & \text{FCbe} &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the above equation, the total classification entries used = 2 and meters used = 2.

9.5.1.12 Example 8

Example

```
sap-ingress 10 create
  num-qos-classifiers 16
  meter 1 create
  exit
  meter 11 multipoint create
  exit
  meter 3 create
  exit
  meter 4 create
  exit
  fc be create
    meter 1
    broadcast-meter 11
    mulitcast-meter 4
  exit
  fc af create
    meter 3
  exit
  default-fc be
  mac-criteria dot1p-only
  entry 1 create
    match dot1p 7 7
    action fc af
  exit
  dot1p 7 fc af
  exit
  match entry 2
    dot1p 5 fc af
  exit
  match entry 3
    dot1p 3 fc af
  exit
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the following number of classification entries per FC:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 & \text{FCh1} &= 0 + 0 + 0 + 0 = 0 & \text{FCef} &= 0 + 0 + 0 + 0 = 0 & \text{FCh2} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCI1} &= 0 + 0 + 0 + 0 = 0 & \text{FCaf} &= 1 + 0 + 1 + 0 = 2 \end{aligned}$$

Because FCaf uses unicast meter, an entry is needed to identify these traffic types explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter "11".

$$\text{FCI2} = 0 + 0 + 0 + 0 = 0 \quad \text{FCbe} = 1 + 1 + 1 + 0 = 3$$

Because FCbe uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

Using the equation, calculate the total classification entries used by this policy, as follows

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (3 * 2)af + (0 * 0)l2 + (1 * 3)be = 9$$

The number of meters used in this policy = 4.

Hence, in this example **num-qos-classifiers 16** is used (that is, maximum of (9, (2 * 4))).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$\begin{aligned} FCnc &= 0 + 0 + 0 + 0 = 0 & FCh1 &= 0 + 0 + 0 + 0 = 0 & FCef &= 0 + 0 + 0 + 0 = 0 & FCh2 &= 0 + 0 + 0 + 0 = 0 \\ FCl1 &= 0 + 0 + 0 + 0 = 0 & FCaf &= 1 + 0 + 0 + 0 = 1 & FCl2 &= 0 + 0 + 0 + 0 = 0 & FCbe &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the equation, calculate the total classification entries used by this policy, as follows:

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (3 * 1)af + (0 * 0)l2 + (1 * 1)be = 4$$

The number of meters used in this policy = 2.

9.5.1.13 Example 9

Example

```
sap-ingress 10 create
  num-qos-classifiers 256
  meter 1 create
  exit
  meter 11 multipoint create
  exit
  meter 3 create
  exit
  meter 4 create
  exit
  fc be create
    meter 1
    broadcast-meter 11
    mulitcast-meter 4
  exit
  fc af create
    meter 3
    broadcast-meter 11
    multicast-meter 4
  exit
  default-fc be
  ip-criteria dscp-only
  entry 1 create
    match dscp cp1
    action fc af
  exit
  entry 2 create
    match dscp cp2
    action fc af
```

```
exit
entry 3 create
    match dscp cp3
    action fc af
exit
entry 4 create
    match dscp cp4
    action fc af
exit
entry 5 create
    match dscp cp5
    action fc af
exit
entry 6 create
    match dscp cp6
    action fc af
exit
entry 7 create
    match dscp cp7
    action fc af
exit
entry 8 create
    match dscp cs1
    action fc af
exit
entry 9 create
    match dscp cp9
    action fc af
exit
entry 10 create
    match dscp af11
    action fc af
exit
entry 11 create
    match dscp cp11
    action fc af
exit
entry 12 create
    match dscp af12
    action fc af
exit
entry 13 create
    match dscp cp13
    action fc af
exit
entry 14 create
    match dscp af13
    action fc af
exit
entry 15 create
    match dscp cp15
    action fc af
exit
entry 16 create
    match dscp cs2
    action fc af
exit
entry 17 create
    match dscp cp17
    action fc af
exit
entry 18 create
    match dscp af21
    action fc af
```



```
exit
entry 19 create
      match dscp cp19
      action fc af
exit
entry 20 create
      match dscp af22
      action fc af
exit
entry 21 create
      match dscp cp21
      action fc af
exit
entry 22 create
      match dscp af23
      action fc af
exit
entry 23 create
      match dscp cp23
      action fc af
exit
entry 24 create
      match dscp cs3
      action fc af
exit
entry 25 create
      match dscp cp25
      action fc af
exit
entry 26 create
      match dscp af31
      action fc af
exit
entry 27 create
      match dscp cp27
      action fc af
exit
entry 28 create
      match dscp af32
      action fc af
exit
entry 29 create
      match dscp cp29
      action fc af
exit
entry 30 create
      match dscp af33
      action fc af
exit
entry 31 create
      match dscp cp31
      action fc af
exit
entry 32 create
      match dscp cs4
      action fc af
exit
entry 33 create
      match dscp cp33
      action fc af
exit
entry 34 create
      match dscp af41
      action fc af
```

```
exit
entry 35 create
      match dscp cp35
      action fc af
exit
entry 36 create
      match dscp af42
      action fc af
exit
entry 37 create
      match dscp cp37
      action fc af
exit
entry 38 create
      match dscp af43
      action fc af
exit
entry 39 create
      match dscp cp39
      action fc af
exit
entry 40 create
      match dscp cs5
      action fc af
exit
entry 41 create
      match dscp cp41
      action fc af
exit
entry 42 create
      match dscp cp42
      action fc af
exit
entry 43 create
      match dscp cp43
      action fc af
exit
entry 44 create
      match dscp cp44
      action fc af
exit
entry 45 create
      match dscp cp45
      action fc af
exit
entry 46 create
      match dscp ef
      action fc af
exit
entry 47 create
      match dscp cp47
      action fc af
exit
entry 48 create
      match dscp ncl
      action fc af
exit
entry 49 create
      match dscp cp49
      action fc af
exit
entry 50 create
      match dscp cp50
      action fc af
```

```

        exit
    exit
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the following number of classification entries per FC:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 & \text{FCh1} &= 0 + 0 + 0 + 0 = 0 & \text{FCef} &= 0 + 0 + 0 + 0 = 0 & \text{FCh2} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCI1} &= 0 + 0 + 0 + 0 = 0 & \text{FCaf} &= 1 + 0 + 1 + 0 = 3 \end{aligned}$$

Because FCaf uses unicast meter, an entry is needed to identify these traffic types explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic to the same FC and use the default meter "11".

$$\text{FCI2} = 0 + 0 + 0 + 0 = 0 \quad \text{FCbe} = 1 + 1 + 1 + 0 = 3$$

Because FCbe uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

Using the equation, calculate the total classification entries used by this policy, as follows:

$$\text{TC} = (0 * 0)\text{nc} + (0 * 0)\text{h1} + (0 * 0)\text{ef} + (0 * 0)\text{h2} + (0 * 0)\text{l1} + (50 * 3)\text{af} + (0 * 0)\text{l2} + (1 * 3)\text{be} = 153$$

The number of meters used in this policy = 4.

Hence, in this example **num-qos-classifiers 256** is used (maximum of (153, (2 * 4)) = 153, rounded off to the next multiple of 2 will be 154).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 & \text{FCh1} &= 0 + 0 + 0 + 0 = 0 & \text{FCef} &= 0 + 0 + 0 + 0 = 0 & \text{FCh2} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCI1} &= 0 + 0 + 0 + 0 = 0 & \text{FCaf} &= 1 + 0 + 0 + 0 = 1 & \text{FCI2} &= 0 + 0 + 0 + 0 = 0 & \text{FCbe} &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the equation, calculate the total classification entries used by this policy, as follows:

$$\text{TC} = (0 * 0)\text{nc} + (0 * 0)\text{h1} + (0 * 0)\text{ef} + (0 * 0)\text{h2} + (0 * 0)\text{l1} + (50 * 1)\text{af} + (0 * 0)\text{l2} + (1 * 1)\text{be} = 51$$

The number of meters used in this policy = 2.

Hence for Epipe SAP it is recommended to define another sap-ingress policy with **num-qos-classifiers 64** is used (that is, maximum of (51, (2 * 2)) = 51, rounded off to the next multiple of 2 will be 52).

9.5.1.14 Example 9a (default multipoint meter "11" is not used):

Example

```

sap-ingress 10 create
    num-qos-classifiers 154
    meter 1 create
    exit
    meter 3 create
    exit

```

```
meter 4 create
exit
meter 11 create
exit

fc be create
    meter 1
    broadcast-meter 11
    multicast-meter 4
exit
fc af create
    meter 3
    broadcast-meter 11
    multicast-meter 4
exit
default-fc be

ip-criteria dscp-only
entry 1 create
    match dscp cp1
    action fc af
exit
entry 2 create
    match dscp cp2
    action fc af
exit
entry 3 create
    match dscp cp3
    action fc af
exit
entry 4 create
    match dscp cp4
    action fc af
exit
entry 5 create
    match dscp cp5
    action fc af
exit
entry 6 create
    match dscp cp6
    action fc af
exit
entry 7 create
    match dscp cp7
    action fc af
exit
entry 8 create
    match dscp cs1
    action fc af
exit
entry 9 create
    match dscp cp9
    action fc af
exit
entry 10 create
    match dscp af11
    action fc af
exit
entry 11 create
    match dscp cp11
    action fc af
exit
entry 12 create
    match dscp af12
```

```
        action fc af
    exit
    entry 13 create
        match dscp cp13
        action fc af
    exit
    entry 14 create
        match dscp af13
        action fc af
    exit
    entry 15 create
        match dscp cp15
        action fc af
    exit
    entry 16 create
        match dscp cs2
        action fc af
    exit
    entry 17 create
        match dscp cp17
        action fc af
    exit
    entry 18 create
        match dscp af21
        action fc af
    exit
    entry 19 create
        match dscp cp19
        action fc af
    exit
    entry 20 create
        match dscp af22
        action fc af
    exit
    entry 21 create
        match dscp cp21
        action fc af
    exit
    entry 22 create
        match dscp af23
        action fc af
    exit
    entry 23 create
        match dscp cp23
        action fc af
    exit
    entry 24 create
        match dscp cs3
        action fc af
    exit
    entry 25 create
        match dscp cp25
        action fc af
    exit
    entry 26 create
        match dscp af31
        action fc af
    exit
    entry 27 create
        match dscp cp27
        action fc af
    exit
    entry 28 create
        match dscp af32
```

```
        action fc af
    exit
    entry 29    create
        match dscp cp29
        action fc af
    exit
    entry 30    create
        match dscp af33
        action fc af
    exit
    entry 31    create
        match dscp cp31
        action fc af
    exit
    entry 32    create
        match dscp cs4
        action fc af
    exit
    entry 33    create
        match dscp cp33
        action fc af
    exit
    entry 34    create
        match dscp af41
        action fc af
    exit
    entry 35    create
        match dscp cp35
        action fc af
    exit
    entry 36    create
        match dscp af42
        action fc af
    exit
    entry 37    create
        match dscp cp37
        action fc af
    exit
    entry 38    create
        match dscp af43
        action fc af
    exit
    entry 39    create
        match dscp cp39
        action fc af
    exit
    entry 40    create
        match dscp cs5
        action fc af
    exit
    entry 41    create
        match dscp cp41
        action fc af
    exit
    entry 42    create
        match dscp cp42
        action fc af
    exit
    entry 43    create
        match dscp cp43
        action fc af
    exit
    entry 44    create
        match dscp cp44
```

```

        action fc af
    exit
    entry 45 create
        match dscp cp45
        action fc af
    exit
    entry 46 create
        match dscp ef
        action fc af
    exit
    entry 47 create
        match dscp cp47
        action fc af
    exit
    entry 48 create
        match dscp nc1
        action fc af
    exit
    entry 49 create
        match dscp cp49
        action fc af
    exit
    entry 50 create
        match dscp cp50
        action fc af
    exit
    exit
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the following number of classification entries per FC:

$$\begin{aligned}
 \text{FCnc} &= 0 + 0 + 0 + 0 = 0 & \text{FCh1} &= 0 + 0 + 0 + 0 = 0 & \text{FCef} &= 0 + 0 + 0 + 0 = 0 & \text{FCh2} &= 0 + 0 + 0 + 0 = 0 \\
 \text{FCI1} &= 0 + 0 + 0 + 0 = 0 & \text{FCaf} &= 1 + 0 + 1 + 0 = 3
 \end{aligned}$$

Because FCaf uses unicast, broadcast and multicast meter, three entries are required to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter. Additionally note that meter "11" is not defined to be multipoint meter, but is used as a normal unicast meter.

$$\text{FCI2} = 0 + 0 + 0 + 0 = 0 \quad \text{FCbe} = 1 + 1 + 1 + 0 = 3$$

Because FCbe uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter. Additionally note that meter "11" is not defined to be multipoint meter, but is used as a normal unicast meter.

Using the equation, calculate the total classification entries used by this policy, as follows:

$$\text{TC} = (0 * 0)\text{nc} + (0 * 0)\text{h1} + (0 * 0)\text{ef} + (0 * 0)\text{h2} + (0 * 0)\text{l1} + (50 * 3)\text{af} + (0 * 0)\text{l2} + (1 * 3)\text{be} = 153$$

The number of meters used in this policy = 4. Hence, in this example **num-qos-classifiers 154** is used (maximum of (153, (2 * 4)) = 153, rounded off to the next multiple of 2 will be 154).

Hence, in this example **num-qos-classifiers 154** is used (maximum of (153, (2 * 4)) = 153, rounded off to the next multiple of 2 will be 154).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$\begin{aligned} FCnc &= 0 + 0 + 0 + 0 = 0 & FCh1 &= 0 + 0 + 0 + 0 = 0 & FCef &= 0 + 0 + 0 + 0 = 0 & FCh2 &= 0 + 0 + 0 + 0 = 0 \\ FCl1 &= 0 + 0 + 0 + 0 = 0 & FCaf &= 1 + 0 + 0 + 0 = 1 & FCl2 &= 0 + 0 + 0 + 0 = 0 & FCbe &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the equation, calculate the total classification entries used by this policy, as follows:

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (50 * 1)af + (0 * 0)l2 + (1 * 1)be = 51$$

The number of meters used in this policy = 2.

Hence for Epipe SAP it is recommended to define another sap-ingress policy with **num-qos-classifiers 52** is used (that is, maximum of (51, (2 * 2)) = 51, rounded off to the multiple of 2 will be 52).

9.5.1.15 Example 10

Example

```
sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 4
  meter 1 create
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
  scope template
  default-fc l2
  fc l2 create
    meter 1
  exit
  fc af create
    meter 1
  exit
  mac-criteria any
    entry 1 create
      match dot1p 7
      action fc af
    exit
  exit
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, compute the number of classification entries per FC as follows:

$$\begin{aligned} FCnc &= 0 + 0 + 0 + 0 = 0 & FCh1 &= 0 + 0 + 0 + 0 = 0 & FCef &= 0 + 0 + 0 + 0 = 0 & FCh2 &= 0 + 0 + 0 + 0 = 0 \\ FCl1 &= 0 + 0 + 0 + 0 = 0 & FCaf &= 1 + 0 + 1 + 0 = 2 & FCl2 &= 1 + 0 + 1 + 0 = 2 & FCbe &= 0 + 0 + 0 + 0 = 2 \end{aligned}$$

Using the equation, calculate the total classification entries used by this policy, as follows:

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (1 * 2)af + (1 * 2)l2 + (0 * 0)be = 4$$

The number of meters used = 2 (because both FCs use meter "1" and meter "11").

Hence, in this example **num-qos-classifiers 4** is used (that is, maximum of (4, (2 * 2))).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$\begin{aligned} FCnc &= 0 + 0 + 0 + 0 = 0 & FCh1 &= 0 + 0 + 0 + 0 = 0 & FCef &= 0 + 0 + 0 + 0 = 0 & FCh2 &= 0 + 0 + 0 + 0 = 0 \\ FCl1 &= 0 + 0 + 0 + 0 = 0 & FCaf &= 1 + 0 + 0 + 0 = 1 & FCl2 &= 1 + 0 + 0 + 0 = 1 & FCbe &= 0 + 0 + 0 + 0 = 0 \end{aligned}$$

Using the above equation, calculate the total classification entries used = 2 and meters used = 1.

As can be seen here, for Epipe SAP with the same amount of resources allocated one can have more FCs if need be.

9.5.1.16 Example 11

Example

```
sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 4
  meter 1 create
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
  scope template
  default-fc be
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, compute the number of classification entries per FC as follows:

$$\begin{aligned} FCnc &= 0 + 0 + 0 + 0 = 0 & FCh1 &= 0 + 0 + 0 + 0 = 0 & FCef &= 0 + 0 + 0 + 0 = 0 & FCh2 &= 0 + 0 + 0 + 0 = 0 \\ FCl1 &= 0 + 0 + 0 + 0 = 0 & FCaf &= 0 + 0 + 0 + 0 = 0 & FCl2 &= 0 + 0 + 0 + 0 = 0 & FCbe &= 1 + 0 + 1 + 0 = 2 \end{aligned}$$

Using the equation, calculate the total classification entries used by this policy, as follows:

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (0 * 0)af + (1 * 2)l2 + (0 * 0)be = 2$$

The number of meters used = 2 (because default FC uses meter "1" and meter "11").

Hence, in this example **num-qos-classifiers 4** is used (that is, maximum of (2, (2 * 2))).

If the same policy were to be used for a SAP in an Epipe service, then because all traffic is classified to a unicast traffic type and because only unicast meters are used, the following:

$$\begin{aligned} FCnc &= 0 + 0 + 0 + 0 = 0 & FCh1 &= 0 + 0 + 0 + 0 = 0 & FCef &= 0 + 0 + 0 + 0 = 0 & FCh2 &= 0 + 0 + 0 + 0 = 0 \\ FCl1 &= 0 + 0 + 0 + 0 = 0 & FCaf &= 0 + 0 + 0 + 0 = 0 & FCl2 &= 0 + 0 + 0 + 0 = 0 & FCbe &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the above equation, total classification entries used = 1 and meters used = 1.

As can be seen here, for Epipe SAP with the same amount of resources allocated one can have more FCs if need be.

9.5.2 Examples: calculating resources required for IP DSCP table-based classification with CAM-based policing (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)

This section provides examples for calculating the amount of resources needed for a service ingress policy when using IP DSCP table-based classification with CAM-based policing. For calculations when CAM-based classification is used, see [Examples: calculating resources required for CAM-based classification](#).

The examples in this section use the two equations below to calculate the value for **num-qos-classifiers** used in the sap-ingress QoS policy. See [Computation of resources used per SAP ingress policy for CAM-based classification](#) for details on these equations.

- total number of classification entries (TC)

$$TC = \sum E(i) * C(i)$$

i = nc, h1, ef, h2, l1, af, l2, be, default-fc

where:

- E(i) is equal to 1 if FC(i) is in use by the dscp-classification policy. Otherwise, E(i) is equal to 0.
- C(i) is the number of classification entries that are required by FC_i to identify different traffic types. For a given FC, for each traffic type (unicast, broadcast, multicast, and unknown-unicast) configured to use a meter, a classification entry is needed.



Note:

The default FC requires one or more additional resources even if all the eight FCs are configured in the dot1p or IP DSCP classification policy, as shown in the example in section [Example 1: Epipe, IES, and VPRN services using unicast traffic type](#).

- total number of QoS resources required (TQ)

$$TQ = \max [(TC), (2 * TP)]$$

where:

- TP is the number of meters/policers used

In addition, the examples show how to determine the number of classification entries for each forwarding class. For example, FCh2 (shown below) is the sum of four traffic types: (unicast (U), broadcast (B), multicast (M), and unknown-unicast (U-u)). See [Calculating the number of classification entries per FC](#) for more information.

$$FCh2 = 1 + 0 + 1 + 0 = 2 \text{ (generalized to } FCh2 = U + B + M + U-u \text{)}$$

If BUM entries are not explicit and multipoint traffic is expected, meter "11" is used and the "M" traffic type is given a "1".

Consider the following items when calculating the resources required when using IP DSCP table-based classification:

- The meters used per FC for different traffic types is in accordance with the rules given in [Service ingress meter selection rules](#) and the number of classification entries per FC is provided in [Determining the number of classification entries](#). In addition:

- If an FC uses one meter for all four traffic types, then the maximum of four classification entries are needed.
 - As a minimum, an FC uses a single meter.
 - Users are provided an option to use between one and four meters per FC, one meter each for four different traffic types, in a VPLS service.
 - Users are provided an option to use either one and two meters per FC, one meter each for two traffic types, in IES and VPRN service when using multicast.
 - Users can use only a single meter per FC, for unicast traffic, in Epipe, IES, and VPRN service, where IES and VPRN are not configured to use multicast.
2. For an Epipe, VPLS, IES, or VPRN SAP, the FCs in use can be determined by counting the FCs configured in the DSCP classification policy and the default FC configured in the SAP ingress policy.
 3. For routed VPLS (RVPLS), to determine the FCs in use, use the following:
 - For SAPs on access port, the FCs in use is the sum total of all FCs configured across the following four items:
 - the DSCP classification policy associated as an override policy under the IP interface context
 - the port DSCP classification policy associated with the access port
 - the untagged-fc value configured under the access port
 - the default-fc configured in the SAP ingress policy
 - For SAPs on hybrid ports, the FCs in use is the sum total of all FCs configured across the following three items:
 - the DSCP classification policy associated as an override policy under the IP interface context
 - the network port policy associated with the hybrid port
 - the default FC value configured under the SAP ingress policy

9.5.2.1 Example 1: Epipe, IES, and VPRN services using unicast traffic type

Example

```
*A:dut-a>config>qos>dscp-classification# info detail
-----
description "dscp-classification-23"
default-dscp-fc "be" profile in
dscp cs3 fc "be"
dscp cs4 fc "ef" profile out
dscp af31 fc "af" profile in
dscp af33 fc "l1" profile in
dscp af41 fc "nc"
dscp cp25 fc "l2"
dscp cp31 fc "h2" profile out
dscp cp33 fc "h1" profile in
-----
*A:dut-a# configure qos sap-ingress 23
```

Example

```
*A:dut-a>config>qos>sap-ingress# info
```

```
-----
description "sap-Ingress-Policy-23"
num-qos-classifiers 16
meter 1 create
    mode trtcm2
    rate cir 5000 pir 7000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 9 create
    mode trtcm2
    rate cir 3000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 10 create
    mode trtcm2
    rate cir 4000 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 11 multipoint create
    mode trtcm2
    rate cir 2000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 12 create
    mode trtcm2
    rate cir 3500 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 13 create
    mode trtcm2
    rate cir 5000 pir 7000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 14 create
    mode trtcm2
    rate cir 5000 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 15 create
    mode trtcm2
    rate cir 4000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
fc "af" create
    meter 11
exit
fc "be" create
    meter 9
exit
fc "ef" create
    meter 14
exit
fc "h1" create
    meter 15
exit
fc "h2" create
```

```

        meter 13
    exit
    fc "l1" create
        meter 12
    exit
    fc "l2" create
        meter 10
    exit
    fc "nc" create
        meter 1
    exit
    dscp-classification 23
-----

```

Example

```

*A:dut-a# configure service epipe 7
*A:dut-a>config>service>epipe# info
-----
description "Default epipe for service id 7"
sap 1/1/3:201 create
    description "Default sap for service id 7"
    ingress
        qos 23  enable-table-classification
    exit
    egress
    exit
exit
no shutdown
-----

```

In the above example, all eight FCs are configured and eight meters are configured. For Epipe service only unicast traffic-type is identified. This requires one classification entry per FC configured and an additional one for the default-fc assignment, for a total of 9 classification entries.

$FC(nc) = 1 + 0 + 0 + 0 = 1$
 $FC(h1) = 1 + 0 + 0 + 0 = 1$
 $FC(ef) = 1 + 0 + 0 + 0 = 1$
 $FC(h2) = 1 + 0 + 0 + 0 = 1$
 $FC(l1) = 1 + 0 + 0 + 0 = 1$
 $FC(af) = 1 + 0 + 0 + 0 = 1$
 $FC(l2) = 1 + 0 + 0 + 0 = 1$
 $FC(be) = 1 + 0 + 0 + 0 = 1$
 $FC(\text{default-fc}) = 1 + 0 + 0 + 0 = 1$

$TC = 1*1 (FC-nc) + 1*1 (FC-h1) + 1*1 (FC-ef) + 1*1 (FC-h2) + 1*1 (FC-l1) + 1*1 (FC-af) + 1*1 (FC-l2) + 1*1 (FC-be) + 1*1 (\text{default-fc}) = 9$

TP = Total meters used is 8.

Hence, in this example **num-qos-classifiers** is set to (maximum (9, (8*2))) = 16.

If the same policy was going to be attached to an IES or VPRN SAP without multicast enabled (VPRN example shown below, IES is similar), the **num-qos-classifiers** required would be the same as that for am Epipe service (that is, 16). The calculations are the same as that for Epipe service.

Example

```

-----
vprn 7 customer 1 create
    description "Default VPRN ID 7"
    route-distinguisher 3:7
    interface "int1VPRN201" create
        address 10.43.44.1/24
        sap 1/1/3:201 create
            description "sap-7-10.43.44.1"

```

```
        ingress
        qos 23 enable-table-classification
        exit
        egress
        exit
    exit
    service-name "XYZ Vprn 7"
    no shutdown
exit
-----
```

9.5.2.2 Example 2: VPLS using unicast and BUM meter with IES or VPRN using multicast

Example

```
*A:dut-a>config>qos>dscp-classification# info detail
-----
description "dscp-classification-23"
default-dscp-fc "be" profile in
dscp cs3 fc "be"
dscp cs4 fc "ef" profile out
dscp af31 fc "af" profile in
dscp af33 fc "l1" profile in
dscp af41 fc "nc"
dscp cp25 fc "l2"
dscp cp31 fc "h2" profile out
dscp cp33 fc "h1" profile in
-----
```

Example

```
*A:dut-a# configure qos sap-ingress 23
*A:dut-a>config>qos>sap-ingress# info
-----
description "sap-Ingress-Policy-23"
num-qos-classifiers 18
meter 1 create
    mode trtcm2
    rate cir 5000 pir 7000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 9 create
    mode trtcm2
    rate cir 3000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 10 create
    mode trtcm2
    rate cir 4000 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 11 multipoint create
    mode trtcm2
    rate cir 2000 pir 5000
    mbs 200 kbits
```

```
        cbs 100 kbits
    exit
meter 12 create
    mode trtcm2
    rate cir 3500 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 13 create
    mode trtcm2
    rate cir 5000 pir 7000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 14 create
    mode trtcm2
    rate cir 5000 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 15 create
    mode trtcm2
    rate cir 4000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
fc "af" create
    meter 11
exit
fc "be" create
    meter 9
exit
fc "ef" create
    meter 14
exit
fc "h1" create
    meter 15
exit
fc "h2" create
    meter 13
exit
fc "l1" create
    meter 12
exit
fc "l2" create
    meter 10
exit
fc "nc" create
    meter 1
exit
dscp-classification 23
-----
```

Example

```
-----
vpls 7 customer 1 svc-sap-type any create
    description "Default tls for service id 7"

sap 1/1/3:201 create
    description "Default sap for service id 7"
    ingress
        qos 23 enable-table-classification
```

```

        exit
    egress
    exit
    exit
    no shutdown
    exit
-----

```

In the above example, all eight FCs are configured and eight meters are configured. In addition, multipoint meter "11" is configured for use. For the VPLS service, four traffic types are identified (unicast and BUM). Because multicast meter "11" is defined, BUM traffic type for all FCs will use meter "11". The number of classification entries required is:

$$\begin{aligned} \text{FC (nc)} &= 1 + 0 + 1 + 0 = 2 & \text{FC (h1)} &= 1 + 0 + 1 + 0 = 2 & \text{FC (ef)} &= 1 + 0 + 1 + 0 = 2 & \text{FC (h2)} &= 1 + 0 + 1 + 0 = 2 \\ \text{FC (l1)} &= 1 + 0 + 1 + 0 = 2 & \text{FC (af)} &= 1 + 0 + 1 + 0 = 2 & \text{FC (l2)} &= 1 + 0 + 1 + 0 = 2 & \text{FC (be)} &= 1 + 0 + 1 + 0 = 2 \\ \text{FC (default-fc)} &= 1 + 0 + 1 + 0 = 2 \end{aligned}$$

$$\text{TC} = 1*2 (\text{FC-nc}) + 1*2 (\text{FC-h1}) + 1*2 (\text{FC-ef}) + 1*2 (\text{FC-h2}) + 1*2 (\text{FC-l1}) + 1*2 (\text{FC-af}) + 1*2 (\text{FC-l2}) + 1*2 (\text{FC-be}) + 1*2 (\text{default-fc}) = 18$$

TP = Total meters used is 8.

Hence, in this example **num-qos-classifiers** is set to (maximum (18, (8*2))) = 18.

If the same policy is attached to an IES or an VPRN service with multicast enabled, then the number of resources required would be 18 (that is, **num-qos-classifiers** needs to be set to 18). The calculations are the same those for VPLS service (shown above).

9.5.2.3 Example 3: VPLS service using unicast, broadcast, multicast, and unknown-unicast with additional FCs

In this example, the sap-ingress policy in Example 2 is changed to define additional meters (shown below) and applied to a SAP configured in a VPLS service.

Example

```

*A:dut-a# configure qos sap-ingress 24
*A:dut-a>config>qos>sap-ingress# info
-----
    description "sap-Ingress-Policy-24"
    num-qos-classifiers 18
    meter 1 create
        mode trtcm2
        rate cir 5000 pir 7000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 9 create
        mode trtcm2
        rate cir 3000 pir 5000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 10 create
        mode trtcm2
        rate cir 4000 pir 6000
        mbs 200 kbits
        cbs 100 kbits

```



```
exit
meter 11 multipoint create
    mode trtcm2
    rate cir 2000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 12 create
    mode trtcm2
    rate cir 3500 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 13 create
    mode trtcm2
    rate cir 5000 pir 7000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 14 create
    mode trtcm2
    rate cir 5000 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 15 create
    mode trtcm2
    rate cir 4000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 20 multipoint create
    mode trtcm2
    rate cir 0 pir 5
    mbs 50 kbits
    cbs 50 kbits
exit
meter 21 multipoint create
    mode trtcm2
    rate cir 0 pir 10
    mbs 100 kbits
    cbs 100 kbits
exit
meter 22 multipoint create
    mode trtcm2
    rate cir 0 pir 50
    mbs 100 kbits
    cbs 100 kbits
exit
fc "af" create
    meter 11
broadcast-meter 22
exit
fc "be" create
    meter 9
exit
fc "ef" create
    meter 14
broadcast-meter 22
exit
fc "h1" create
    meter 15
broadcast-meter 22
exit
```

```
fc "h2" create
    meter 13
broadcast-meter 22
exit
fc "l1" create
    meter 12
broadcast-meter 22
    unknown-unicast 21
exit
fc "l2" create
    meter 10
broadcast-meter 20
exit
fc "nc" create
    meter 1
exit
dscp-classification 23
-----
```

For the VPLS service, four traffic types are identified (unicast and BUM). Because multicast meter "11" is defined, BUM traffic types for all FCs will use meter "11" (since the user has not configured an explicit multicast-meter for the FC). Hence, the number of classification entries required is determined as shown below:

FC (nc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (h1) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (ef) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (h2) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (l1) = 1 + 1 + 1 + 1 = 4 (one for unicast, one for broadcast, one for unknown-unicast & one for multicast) FC (af) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast and unknown-unicast) FC (l2) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast and unknown-unicast) FC (be) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (default-fc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM, since default-dscp-fc is configured to be FC 'be' in the dscp-classification policy)

$$TC = 1*2 (FC-nc) + 1*3 (FC-h1) + 1*3 (FC-ef) + 1*3 (FC-h2) + 1*4 (FC-l1) + 1*3 (FC-af) + 1*3 (FC-l2) + 1*2 (FC-be) + 1*2 (default-fc) = 25$$

TP = Total meters used is 11.

Hence, in this example, **num-qos-classifiers** is calculated to be maximum (25, (11*2)) = 25, but is set to 26 after rounding off to the next highest even number.

If the same policy is attached to an IES or an VPRN service with multicast enabled, then the number of resources required would be 18 (that is, **num-qos-classifiers** needs to be set to 18). This is because for IES and VPRN only unicast and multicast traffic types are supported; broadcast and unknown-unicast traffic types are not supported and do not consume any resources. The calculation is as shown below:

FC (nc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (h1) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (ef) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (h2) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (l1) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (af) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (l2) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (be) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (default-fc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM, since default-dscp-fc is configured to be FC 'be' in the dscp-classification policy)

$$TC = 1*2 (FC-nc) + 1*2 (FC-h1) + 1*2 (FC-ef) + 1*2 (FC-h2) + 1*2 (FC-l1) + 1*2 (FC-af) + 1*2 (FC-l2) + 1*2 (FC-be) + 1*2 (default-fc) = 18$$

TP = Total meters used is 8 (meter 20, meter 21, and meter 22 are not used because they are associated with broadcast and unknown-unicast traffic types, which are not supported for IES and VPRN).

Hence, in this example **num-qos-classifiers** is set to (maximum (18, (8*2)) = 18.

9.5.2.4 Example 4: routed VPLS on access port using unicast, broadcast, multicast, and unknown-unicast with additional FCs

Example

```
*A:dut-a>config>qos>dscp-classification# info detail
-----
description "dscp-classification-23"
default-dscp-fc "be" profile in
dscp cs3 fc "be"
dscp cs4 fc "ef" profile out
dscp af31 fc "af" profile in
dscp af33 fc "l1" profile in
dscp af41 fc "nc"
dscp cp25 fc "l2"
dscp cp31 fc "h2" profile out
dscp cp33 fc "h1" profile in
-----
```

Example

```
*A:dut-a# configure qos sap-ingress 24
*A:dut-a>config>qos>sap-ingress# info
-----
description "sap-Ingress-Policy-24"
num-qos-classifiers 26
meter 1 create
mode trtcm2
rate cir 5000 pir 7000
mbs 200 kbits
cbs 100 kbits
exit
meter 9 create
mode trtcm2
rate cir 3000 pir 5000
mbs 200 kbits
cbs 100 kbits
exit
meter 10 create
mode trtcm2
rate cir 4000 pir 6000
mbs 200 kbits
cbs 100 kbits
exit
meter 11 multipoint create
mode trtcm2
rate cir 2000 pir 5000
mbs 200 kbits
cbs 100 kbits
exit
meter 12 create
mode trtcm2
rate cir 3500 pir 6000
mbs 200 kbits
```

```
        cbs 100 kbits
    exit
    meter 13 create
        mode trtcm2
        rate cir 5000 pir 7000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 14 create
        mode trtcm2
        rate cir 5000 pir 6000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 15 create
        mode trtcm2
        rate cir 4000 pir 5000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 20 multipoint create
        mode trtcm2
        rate cir 0 pir 5
        mbs 50 kbits
        cbs 50 kbits
    exit
    meter 21 multipoint create
        mode trtcm2
        rate cir 0 pir 10
        mbs 100 kbits
        cbs 100 kbits
    exit
    meter 22 multipoint create
        mode trtcm2
        rate cir 0 pir 50
        mbs 100 kbits
        cbs 100 kbits
    exit
    fc "af" create
        meter 11
    broadcast-meter 22
    exit
    fc "be" create
        meter 9
    exit
    fc "ef" create
        meter 14
    broadcast-meter 22
    exit
    fc "h1" create
        meter 15
    broadcast-meter 22
    exit
    fc "h2" create
        meter 13
    broadcast-meter 22
    exit
    fc "l1" create
        meter 12
    broadcast-meter 22
        unknown-unicast 21
    exit
    fc "l2" create
        meter 10
```

```

broadcast-meter 20
exit
fc "nc" create
    meter 1
exit
dscp-classification 23
-----

```

IES service and routed VPLS service configuration is shown below.

Example

```

-----
ies 7 customer 1 vpn 7 create
    description "Default Ies service id 7"
    interface "int1IES201" create
        address 10.43.44.1/24
        vpls "rvpls607"
        ingress
            enable-table-classification
            routed-override-qos-policy 23
        exit
    exit
exit
service-name "XYZ Ies 7"
no shutdown
-----

```

Example

```

-----
vpls 607 customer 1 r-vpls svc-sap-type any create
    description "Default tls service id 607"
    allow-ip-int-bind
    exit
    stp
        shutdown
    exit
    service-name "rvpls607"
    sap 1/1/3:201 create
        description "Default sap service id 607"
        ingress
            qos 23 enable-table-classification
        exit
        egress
        exit
    exit
no shutdown
-----

```

The configuration for the DSCP classification policy associated with the access port that the RVPLS SAP is configured on, and which is used for classifying bridged packets is shown below:

Example

```

*A:dut-a# configure port 1/1/3
*A:dut-a>config>port# info
-----
    ethernet
        mode access
        enable-table-classification

```

```

access
  ingress
    dscp-classification 23
    untagged-fc ef
  exit
exit
encap-type dot1q
mtu 9212
exit
no shutdown
-----

```

For the bridged traffic in a VPLS service four traffic types are identified (unicast and BUM). Because multicast meter "11" is defined, BUM traffic type for all FCs will use meter "11". (An explicit multicast-meter for the FC has not been configured).

The number of classification entries required is as follows:

FC (nc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (h1) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (ef) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (h2) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (l1) = 1 + 1 + 1 + 1 = 4 (one for unicast, one for broadcast, one for unknown-unicast & one for multicast) FC (af) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast and unknown-unicast) FC (l2) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast and unknown-unicast) FC (be) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (default-fc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM, because default-dscp-fc is configured to be FC 'be' in the dscp-classification policy)

$$TC = 1*2 (FC-nc) + 1*3 (FC-h1) + 1*3 (FC-ef) + 1*3 (FC-h2) + 1*4 (FC-l1) + 1*3 (FC-af) + 1*3 (FC-l2) + 1*2 (FC-be) + 1*2 (default-fc) = 25$$

TP = Total meters used is 11.

Hence, in this example **num-qos-classifiers** is set to (maximum (25, (11*2)) = 25, which means 26 after rounding off to the next highest even number.



Note:

For routed traffic in the routed VPLS service only the unicast traffic type is supported currently. This does not change the amount of resources needed since bridged traffic requires higher amount of resources. To reduce the amount of resources, users can dedicate a single meter for BUM traffic from all FCs, as shown in the following example (note that meter "11" is used for all FCs automatically when meter "11" is defined in the policy).

Example

```

*A:dut-a# configure qos sap-ingress 34
*A:dut-a>config>qos>sap-ingress# info
-----
description "sap-Ingress-Policy-34"
num-qos-classifiers 18
meter 1 create
  mode trtcm2
  rate cir 5000 pir 7000
  mbs 200 kbits
  cbs 100 kbits
exit
meter 9 create
  mode trtcm2

```

```
        rate cir 3000 pir 5000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 10 create
        mode trtcm2
        rate cir 4000 pir 6000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 11 multipoint create
        mode trtcm2
        rate cir 2000 pir 5000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 12 create
        mode trtcm2
        rate cir 3500 pir 6000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 13 create
        mode trtcm2
        rate cir 5000 pir 7000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 14 create
        mode trtcm2
        rate cir 5000 pir 6000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 15 create
        mode trtcm2
        rate cir 4000 pir 5000
        mbs 200 kbits
        cbs 100 kbits
    exit
    meter 20 multipoint create
        mode trtcm2
        rate cir 400000 pir 1000000
        mbs 5000 kbits
        cbs 500 kbits
    exit
    fc "af" create
        meter 9
    exit
    fc "be" create
        meter 20
    exit
    fc "ef" create
        meter 14
    exit
    fc "h1" create
        meter 15
    exit
    fc "h2" create
        meter 13
    exit
    fc "l1" create
        meter 12
    exit
```

```
fc "l2" create
    meter 10
exit
fc "nc" create
    meter 1
exit
dscp-classification 23
-----
```

In the above example, all eight FCs are configured and nine meters are configured, with multipoint meter "11" dedicated to all BUM traffic for all FCs in use. For the VPLS service four traffic-types are identified (unicast and BUM). Because multicast meter "11" is defined, BUM traffic type for all FCs will use meter "11". The number of classification entries required is as follows:

FC (nc) = 1 + 0 + 1 + 0 = 2 FC (h1) = 1 + 0 + 1 + 0 = 2 FC (ef) = 1 + 0 + 1 + 0 = 2 FC (h2) = 1 + 0 + 1 + 0 = 2 FC (l1) = 1 + 0 + 1 + 0 = 2 FC (af) = 1 + 0 + 1 + 0 = 2 FC (l2) = 1 + 0 + 1 + 0 = 2 FC (be) = 1 + 0 + 1 + 0 = 2 FC (default-fc) = 1 + 0 + 1 + 0 = 2

TC = 1*2 (FC-nc) + 1*2 (FC-h1) + 1*2 (FC-ef) + 1*2 (FC-h2) + 1*2 (FC-l1) + 1*2 (FC-af) + 1*2 (FC-l2) + 1*2 (FC-be) + 1*2 (default-fc) = 18

TP = Total meters used is 9.

Hence, in this example **num-qos-classifiers** is set to (maximum (18, (9*2)) = 18.

9.5.2.5 Example 5: routed VPLS service on a hybrid port using unicast, broadcast, multicast and unknown-unicast for some FCs

Example

```
*A:dut-a>config>qos>dscp-classification# info detail
-----
description "dscp-classification-23"
default-dscp-fc "be" profile in
dscp cs3 fc "be"
dscp cs4 fc "ef" profile out
dscp af31 fc "af" profile in
dscp af33 fc "l1" profile in
dscp af41 fc "nc"
dscp cp25 fc "l2"
dscp cp31 fc "h2" profile out
dscp cp33 fc "h1" profile in
-----
```

Example

```
*A:dut-a# configure qos sap-ingress 24
*A:dut-a>config>qos>sap-ingress# info
-----
description "sap-Ingress-Policy-24"
num-qos-classifiers 26
meter 1 create
    mode trtcm2
    rate cir 5000 pir 7000
    mbs 200 kbits
    cbs 100 kbits
-----
```



```
exit
meter 9 create
    mode trtcm2
    rate cir 3000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 10 create
    mode trtcm2
    rate cir 4000 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 11 multipoint create
    mode trtcm2
    rate cir 2000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 12 create
    mode trtcm2
    rate cir 3500 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 13 create
    mode trtcm2
    rate cir 5000 pir 7000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 14 create
    mode trtcm2
    rate cir 5000 pir 6000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 15 create
    mode trtcm2
    rate cir 4000 pir 5000
    mbs 200 kbits
    cbs 100 kbits
exit
meter 20 multipoint create
    mode trtcm2
    rate cir 0 pir 5
    mbs 50 kbits
    cbs 50 kbits
exit
meter 21 multipoint create
    mode trtcm2
    rate cir 0 pir 10
    mbs 100 kbits
    cbs 100 kbits
exit
meter 22 multipoint create
    mode trtcm2
    rate cir 0 pir 50
    mbs 100 kbits
    cbs 100 kbits
exit
fc "af" create
    meter 11
    broadcast-meter 22
```

```
exit
fc "be" create
    meter 9
exit
fc "ef" create
    meter 14
    broadcast-meter 22
exit
fc "h1" create
    meter 15
    broadcast-meter 22
exit
fc "h2" create
    meter 13
    broadcast-meter 22
exit
fc "l1" create
    meter 12
    broadcast-meter 22
    unknown-unicast 21
exit
fc "l2" create
    meter 10
    broadcast-meter 20
exit
fc "nc" create
    meter 1
exit
dscp-classification 23
-----
```

The IES service and routed VPLS service configuration is as follows:

Example

```
-----
ies 7 customer 1 vpn 7 create
    description "Default Ies service id 7"
    interface "int1IES201" create
        address 10.43.44.1/24
        vpls "rvpls607"
            ingress
                enable-table-classification
                routed-override-qos-policy 23
            exit
        exit
    exit
    service-name "XYZ Ies 7"
    no shutdown
-----
```

Example

```
-----
vpls 607 customer 1 r-vpls svc-sap-type any create
    description "Default tls service id 607"
    allow-ip-int-bind
    exit
    stp
        shutdown
    exit
    service-name "rvpls607"
-----
```

```
sap 1/1/3:201 create
description "Default sap service id 607"
ingress
    qos 24 enable-table-classification
exit
egress
exit
exit
no shutdown
-----
```

The configuration for the network port policy for a hybrid port, followed by associating the policy with hybrid port 1/1/3 is shown below. This configuration is used for classifying RVPLS SAP bridged packets and also for classifying IP traffic received and processed in the context of the network port IP interface.

Example

```
**A:dut-a >config>qos# network 23
**A:dut-a >config>qos>network# info
-----
    ingress
        default-action fc be profile in
        meter 1 create
    exit
    dscp cs4 fc ef profile out
    dscp af31 fc af profile in
    dscp af33 fc l1 profile in
    dscp af41 fc nc profile in
    dscp cp25 fc l2 profile out
    dscp cp31 fc h2 profile out
    dscp cp33 fc h1 profile in
    exit
    egress
    exit
-----
```

Example

```
*A:dut-a #/configure port 1/1/3
A:NS1543C2102>config>port# info
-----
    shutdown
    ethernet
        mode hybrid
        access
    exit
    encap-type dot1q
    network
        queue-policy "_tmnx_hybrid_default"
        qos 23
    exit
    exit
-----
*A:NS1543C2102>config>port#
```

For the bridged traffic in a VPLS four traffic types are identified (unicast and BUM). Because multicast meter "11" is defined, BUM traffic type for all FCs will use meter "11". (The user has not configured an explicit multicast-meter for the FC). The number of classification entries required is as follows.

FC (nc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (h1) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (ef) = 1 + 1 + 1 + 0 = 3 (one

for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (h2) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (l1) = 1 + 1 + 1 + 1 = 4 (one for unicast, one for broadcast, one for unknown-unicast & one for multicast) FC (af) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast and unknown-unicast) FC (l2) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast and unknown-unicast) FC (be) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (default-fc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM, since default-dscp-fc is configured to be FC 'be' in the dscp-classification policy)

$$TC = 1*2 (FC-nc) + 1*3 (FC-h1) + 1*3 (FC-ef) + 1*3 (FC-h2) + 1*4 (FC-l1) + 1*3 (FC-af) + 1*3 (FC-l2) + 1*2 (FC-be) + 1*2 (default-fc) = 25$$

TP = Total meters used is 11.

Hence, in this example **num-qos-classifiers** is calculated to be (maximum (25, (11*2)) = 25, but is set to 26 after rounding off to the next highest even number.

9.5.2.6 Example 6: routed VPLS on access port and hybrid port

Example 6 is similar to Example 5, except that FCs in use from the port policy and the override policy are considered. The following configuration shows the policy association with two SAPs configured in the routed VPLS, with one SAP on the access port and the other SAP on the hybrid port.

Example

```
*A:dut-a>config>qos>dscp-classification 23 create
*A:dut-a>config>qos>dscp-classification# info detail
-----
description "dscp-classification-23"
default-dscp-fc "be" profile in
dscp cs3 fc "be"
dscp cs4 fc "ef" profile out
dscp af31 fc "af" profile in
dscp af33 fc "l1" profile in
dscp af41 fc "nc"
dscp cp25 fc "l2"
dscp cp31 fc "h2" profile out
dscp cp33 fc "h1" profile in
-----
```

Example

```
*A:dut-a>config>qos>dscp-classification 24 create
*A:dut-a>config>qos>dscp-classification# info detail
-----
description "dscp-classification-24"
default-dscp-fc "be" profile in
dscp cs4 fc "ef" profile out
dscp af31 fc "af" profile in
-----
```

Example

```
*A:dut-a# configure qos sap-ingress 24
*A:dut-a>config>qos>sap-ingress# info
-----
description "sap-Ingress-Policy-24"
```

```
num-qos-classifiers 26
meter 1 create
  mode trtcm2
  rate cir 5000 pir 7000
  mbs 200 kbits
  cbs 100 kbits
exit
meter 9 create
  mode trtcm2
  rate cir 3000 pir 5000
  mbs 200 kbits
  cbs 100 kbits
exit
meter 10 create
  mode trtcm2
  rate cir 4000 pir 6000
  mbs 200 kbits
  cbs 100 kbits
exit
meter 11 multipoint create
  mode trtcm2
  rate cir 2000 pir 5000
  mbs 200 kbits
  cbs 100 kbits
exit
meter 12 create
  mode trtcm2
  rate cir 3500 pir 6000
  mbs 200 kbits
  cbs 100 kbits
exit
meter 13 create
  mode trtcm2
  rate cir 5000 pir 7000
  mbs 200 kbits
  cbs 100 kbits
exit
meter 14 create
  mode trtcm2
  rate cir 5000 pir 6000
  mbs 200 kbits
  cbs 100 kbits
exit
meter 15 create
  mode trtcm2
  rate cir 4000 pir 5000
  mbs 200 kbits
  cbs 100 kbits
exit
meter 20 multipoint create
  mode trtcm2
  rate cir 0 pir 5
  mbs 50 kbits
  cbs 50 kbits
exit
meter 21 multipoint create
  mode trtcm2
  rate cir 0 pir 10
  mbs 100 kbits
  cbs 100 kbits
exit
meter 22 multipoint create
  mode trtcm2
  rate cir 0 pir 50
```

```
        mbs 100 kbits
        cbs 100 kbits
    exit
    fc "af" create
        meter 11
        broadcast-meter 22
    exit
    fc "be" create
        meter 9
    exit
    fc "ef" create
        meter 14
        broadcast-meter 22
    exit
    fc "h1" create
        meter 15
        broadcast-meter 22
    exit
    fc "h2" create
        meter 13
        broadcast-meter 22
    exit
    fc "l1" create
        meter 12
        broadcast-meter 22
        unknown-unicast 21
    exit
    fc "l2" create
        meter 10
        broadcast-meter 20
    exit
    fc "nc" create
        meter 1
    exit
    dscp-classification 24
-----
```

The following are sample IES service and routed VPLS service configuration outputs.

Example

```
-----
ies 7 customer 1 vpn 7 create
    description "Default Ies service id 7"
    interface "int1IES201" create
        address 10.43.44.1/24
        vpls "rvpls607"
            ingress
                enable-table-classification
                routed-override-qos-policy 24
            exit
        exit
    exit
    service-name "XYZ Ies 7"
    no shutdown
-----
```

Example

```
-----
vpls 607 customer 1 r-vpls svc-sap-type any create
    description "Default tls service id 607"
```

```

allow-ip-int-bind
exit
stp
    shutdown
exit
service-name "rvpls607"
sap 1/1/24:201 create // SAP on hybrid port
    description "Default sap service id 607"
    ingress
        qos 24 enable-table-classification
    exit
    egress
    exit
exit
sap 1/1/3:201 create // SAP on access port
    description "Default sap service id 607"
    ingress
        qos 24 enable-table-classification
    exit
    egress
    exit
exit
no shutdown
-----

```

The configuration of the network port policy for a hybrid port, followed by associating the policy with hybrid port 1/1/3 is shown below. This configuration is used for classifying RVPLS SAP bridged packets and also for classifying IP traffic received and processed in the context of the network port IP interface.

Example

```

**A:dut-a >config>qos# network 123
**A:dut-a >config>qos>network# info
-----
    ingress
        default-action fc be profile in
        meter 1 create
        exit
        dscp cs4 fc ef profile out
        dscp af31 fc af profile in
        dscp af33 fc l1 profile in
        dscp af41 fc nc profile in
        dscp cp25 fc l2 profile out
        dscp cp31 fc h2 profile out
        dscp cp33 fc h1 profile in
    exit
    egress
    exit
-----

```

Example

```

*A:dut-a #/configure port 1/1/24
A:NS1543C2102>config>port# info
-----
    shutdown
    ethernet
        mode hybrid
        access
    exit
    encap-type dot1q
    network

```

```

        queue-policy "_tmnx_hybrid_default"
        qos 123
    exit
exit
-----
*A:NS1543C2102>config>port#

```

The configuration of the DSCP classification policy associated with the access port, which is the port that the RVPLS SAP is configured on, is shown below. The policy is used for classifying bridged packets.

Example

```

*A:dut-a# configure port 1/1/3
*A:dut-a>config>port# info
-----
    ethernet
        mode access
        enable-table-classification
        access
            ingress
                dscp-classification 24
                untagged-fc ef
            exit
        exit
        encap-type dot1q
        mtu 9212
    exit
    no shutdown
-----

```

To determine the resources needed for RVPLS SAP 1/1/3:201 consider the FCs configured in the DSCP classification policy configured in access port 1/1/3 context—there are three FCs (be, af, ef) configured. In addition, consider the FCs configured in the DSCP classification policy configured in the context of IES IP interface—there are three FCs (be, af, ef) configured. The total number of resources for RVPLS SAP 1/1/3:201 is computed as follows using the meter configuration under SAP ingress policy 24.

FC (ef) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (af) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast and unknown-unicast) FC (be) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (default-fc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM)

$$TC = 0*2 (FC-nc) + 0*3 (FC-h1) + 1*3 (FC-ef) + 0*3 (FC-h2) + 0*4 (FC-l1) + 1*3 (FC-af) + 0*3 (FC-l2) + 1*2 (FC-be) + 1*2 (default-fc) = 10$$

TP = Total meters used is 4 (meter 11, meter 9, meter 22, and meter 14).

Hence, to use the available resources efficiently for the RVPLS SAP 1/1/3:201 we need a policy with **num-qos-classifiers** set to (maximum (10, (4*2))) = 10.

To determine the resources needed for RVPLS SAP 1/1/24:201 consider the FCs configured in the network port policy 123, which is associated with network port 1/1/24—all eight FCs are configured. In addition, consider the FCs configured in the DSCP classification policy configured in the context of IES IP interface—three FCs (be, af, ef) are configured. The total number of resources for RVPLS SAP 1/1/24:201 is computed as follows, using the meter configuration under the SAP ingress policy 24:

FC (nc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (h1) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (ef) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (h2) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast & unknown-unicast) FC (l1) = 1 + 1 + 1 + 1 =

4 (one for unicast, one for broadcast, one for unknown-unicast & one for multicast) FC (af) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast and unknown-unicast) FC (l2) = 1 + 1 + 1 + 0 = 3 (one for unicast, one for broadcast, and one for both multicast and unknown-unicast) FC (be) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM) FC (default-fc) = 1 + 0 + 1 + 0 = 2 (one for unicast and one for all of BUM, because default-dscp-fc is configured to be FC 'be' in the dscp-classification policy)

$$TC = 1*2 (FC-nc) + 1*3 (FC-h1) + 1*3 (FC-ef) + 1*3 (FC-h2) + 1*4 (FC-l1) + 1*3 (FC-af) + 1*3 (FC-l2) + 1*2 (FC-be) + 1*2 (default-fc) = 25$$

TP = Total meters used is 11.

Hence, for RVPLS SAP /1/124:201, we need a policy with **num-qos-classifiers** in the policy is calculated to be (maximum (25, (11*2)) = 25, but is set to 26 after rounding off to the next highest even number.

9.6 Basic configurations

A basic service ingress QoS policy must conform to the following:

- Have a unique service ingress QoS policy ID.
- Allocate number of classifier and meter resources needed for use
- Have a QoS policy scope of **template** or **exclusive**.
- Have at least one default unicast forwarding class meter/queue.
- (optionally) Use multipoint forwarding class meter/queues.

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP, a default QoS policy is applied.

9.6.1 Create service ingress QoS policies

To create a service ingress QoS policy, define the following:

- A policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Specify the **num-qos-classifiers** *num-resources* parameter. The default value is 2.
- Specify a default forwarding class for the policy. All packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class.
- Define forwarding class parameters:
 - Modify the **unicast-meter/queue** default value to override the default unicast forwarding type meter mapping for **fc fc-name**.
 - Modify the **multicast-meter/queue** default value to override the default multicast forwarding type meters/queue mapping for **fc fc-name**.
 - Modify the **unknown-meter/queue** default value to override the default unknown unicast forwarding type meter mapping for **fc fc-name**.
 - Modify the **broadcast-meter** default value to override the default broadcast forwarding type meter mapping for **fc fc-name**.

- On platforms where applicable, specify the appropriate classification criteria: IPv4/IPv6 or MAC criteria, or both IP and MAC criteria. You can define IPv4/IPv6, MAC-based, and MAC- and IP-based SAP ingress policies to select the appropriate ingress meter and corresponding forwarding class for matched traffic. See [Service ingress IP match criteria](#) and [Service ingress MAC match criteria](#).
- A SAP ingress policy is created with a **template** scope. The scope can be modified to **exclusive** for a special one-time use policy. Otherwise, the **template** scope enables the policy to be applied to multiple SAPs.

Output example

The following is a sample service ingress policy configuration output.

```
A:ALA-7>config>qos>sap-ingress# info
-----
...
    sap-ingress 100 create
      description "Used on VPN sap"
      num-qos-classifications 2
      no meter
      dscp-classification 100      (7210 SAS-Mxp only)
      no ip-criteria
      no ipv6-criteria
      no mac-criteria
      scope template
...
-----
A:ALA-7>config>qos>sap-ingress#
```

9.6.1.1 Service ingress QoS meter

To create service ingress meter parameters, define the following:

- A new meter ID value — The system will not dynamically assign a value.
- Meter parameters — Ingress meters support the definition of either srTCM (Single Rate Tri-Color Meter) or trTCM (Two Rate Tri-Color Meter), CIR/PIR, CBS/MBS parameters.

Output example:

The following is a sample ingress meter configuration output.

```
A:ALA-7>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
...
sap-ingress 100 create
  description "Used on VPN sap"
  meter 1 create
  exit
  meter 11 multipoint create
  exit
  meter 2 create
    rate cir 11000
  exit
  meter 3 create
    cbs 32
    rate 11000
  exit
```

```
meter 4 create
  rate 1
exit
meter 5 create
  cbs 64
  mbs 128
  rate cir 1500 pir 1500
exit
meter 6 create
  mode srtcm
  rate cir 2500 pir 2500
exit
meter 7 create
  cbs 256
  mbs 512
  rate cir 100 pir 36
exit
meter 8 create
  cbs 256
  mbs 512
  rate cir 11000
exit
meter 9 create
  rate cir 11000
exit
meter 10 create
  rate cir 1
exit
meter 12 create
  rate cir 1500 pir 1500
exit
meter 13 create
  rate cir 2500 pir 2500
exit
meter 14 create
  rate cir 36 pir 100
exit
meter 15 create
  rate cir 36 pir 100
exit
meter 16 create
  cbs 128
  mbs 256
  rate cir 36 pir 100
exit
...
#-----
A:ALA-7>config>qos#
```

9.6.1.2 Service ingress IP match criteria

When specifying SAP ingress match criteria, only one match criteria type can be configured in the SAP ingress QoS policy.

Output example

The following are two sample ingress IP criteria configuration outputs.

```
7210-SAS>config>qos>sap-ingress# info
-----
num-qos-classifiers 32
```

```
meter 1 create
exit
meter 11 multipoint create
exit
fc "h2" create
exit
ip-criteria any
  entry 16 create
    description "test"
    match
    exit
    action fc "be"
  exit
exit
-----
7210-SAS>config>qos>sap-ingress#

7210-SAS>config>qos>sap-ingress# info
-----
num-qos-classifiers 4
meter 1 create
exit
meter 11 multipoint create
exit
ip-criteria dscp-only
  entry 30 create
    match
    exit
    action fc "l2"
  exit
exit
-----
7210-SAS>config>qos>sap-ingress#
```

9.6.1.3 Service ingress MAC match criteria

To configure service ingress QoS 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.
- **action**
The action to associate the forwarding class with a specific MAC criteria entry ID.
- **a description**
The description provides a brief overview of policy features.

Output example:

The following is a sample ingress MAC criteria configuration output.

```
7210-SAS>config>qos>sap-ingress# info
-----
description "test"
num-qos-classifiers 16
meter 1 create
exit
meter 11 multipoint create
exit
```

```
        mac-criteria dot1p-only
        entry 25 create
            match
            exit
            no action
        exit
    exit
    default-fc "h1"
-----
7210-SAS>config>qos>sap-ingress#
```

9.6.2 Applying service ingress policies

This section describes applying SAP ingress policies to service SAPs.

9.6.2.1 Epipe

The following is a sample Epipe service configuration output with SAP ingress policy 100 applied to the SAP. The **enable-table-classification** keyword applies only to the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

Output example

```
A:ALA-7>config>service# info
-----
    epipe 6 customer 6 vpn 6 create
        description "Epipe service to west coast"
        sap 1/1/10:10 create
            exit
            ingress
                qos 100 [enable-table-classification]
            exit
        exit
    exit
-----
A:ALA-7>config>service#
```

9.6.2.2 VPLS

The following is a sample VPLS service configuration output with SAP ingress policy 100. The **enable-table-classification** keyword applies only to the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

Output example:

```
A:ALA-7>config>service# info
-----
    vpls 700 customer 7 vpn 700 create
        description "test"
        stp
            shutdown
        exit
        sap 1/1/9:10 create
            ingress
                qos 100 [enable-table-classification]
            exit
    exit
-----
```

```
        exit
    exit
-----
A:ALA-7>config>service#
```

9.6.2.3 VPRN

The following is a sample VPRN service configuration output. The **enable-table-classification** keyword applies only to the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

Output example:

```
A:ALA-7>config>service# info
-----
...
    vprn 1 customer 1 create
        autonomous-system 10000
        route-distinguisher 10001:1
        auto-bind 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 [enable-table-classification]
                exit
            exit
        exit
    exit
    no shutdown
exit
...
-----
A:ALA-7>config>service#
```

9.6.2.4 IES

The following is a sample IES service configuration output. The **enable-table-classification** keyword applies only to the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

Output example

```
A:ALA-7>config>service# info
-----
...
    ies 1 customer 1 create
        interface "to-cl" create
            address 10.1.0.1/24
            sap 1/1/10:100 create
                ingress
                    qos 100 [enable-table-classification]
                exit
            exit
        exit
    exit
    no shutdown
exit
...
-----
```

```
A:ALA-7>config>service#
```

9.7 Service management tasks

This section describes service management tasks.

9.7.1 Deleting QoS policies

Every service SAP is associated, by default, with the appropriate ingress policy (*policy-id* 1). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the SAP configuration. When you remove a non-default service ingress policy, the association reverts to the default *policy-id* 1.

A QoS policy cannot be deleted until it is removed from all SAPs where they are applied.

```
A:ALA-7>config>qos# no sap-ingress 100
MINOR: CLI SAP ingress policy "100" cannot be removed because it is in use.
A:ALA-7>config>qos#
```

9.7.1.1 Remove a QoS policy from service SAPs

The following Epipe service output examples show that the SAP service ingress reverted to *policy-id* "1" when the non-default policies were removed from the configuration.

Output example

```
A:ALA-104>config>service>epipe# info detail
-----
      description "Distributed Epipe service to west coast"
      no tod-suite
      dotlag
      exit
      ingress
        qos 1
        no filter
      exit
      egress
        no filter
      exit
      no collect-stats
      no accounting-policy
      no shutdown
-----
A:ALA-7>config>service>epipe#
```

9.7.2 Copying and overwriting QoS policies

You can copy an existing service ingress policy, rename it with a new policy ID value, or overwrite an existing policy ID. The overwrite option must be specified or an error occurs if the destination policy ID exists.

Use the following syntax to copy and overwrite QoS policies.

```
config>qos# copy {sap-ingress} source-policy-id dest-policy-id [overwrite]
```

Example

```
*A:ALU-7210>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
    sap-ingress 100 create
        description "Used on VPN sap"
        meter 1 create
        exit
        meter 2 multipoint create
        exit
        meter 10 create
            rate cir 11000
        exit
        meter 11 multipoint create
        exit
    exit
    sap-ingress 101 create
        description "Used on VPN sap"
        meter 1 create
        exit
        meter 2 multipoint create
        exit
        meter 10 create
            rate cir 11000
        exit
        meter 11 multipoint create
        exit
    exit
    sap-ingress 200 create
        description "Used on VPN sap"
        meter 1 create
        exit
        meter 2 multipoint create
        exit
        meter 10 create
            rate cir 11000
        exit
        meter 11 multipoint create
        exit
    exit
-----
*A:ALU-7210>config>qos#
```

9.7.3 Remove a policy from the QoS configuration

Use the following syntax to remove a policy from the QoS configuration.

```
config>qos# no sap-ingress policy-id
```

Example:

```
config>qos# no sap-ingress 100
```


9.7.4 Editing QoS policies

You can change QoS existing policies and entries. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors copy the policy to a work area, make the edits, and then write over the original policy.

9.8 Service ingress QoS policy command reference

9.8.1 Command hierarchy

- [Service ingress QoS policy commands](#)
- [Table-based IP DSCP and dot1p classification policy commands for SAP ingress \(7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12\)](#)
- [Service meter commands for SAP ingress \(7210 SAS-Mxp only\)](#)
- [Operational commands](#)
- [Show commands](#)

9.8.1.1 Service ingress QoS policy commands

```
config
- qos
- [no] sap-ingress policy-id [create] [use-svc-meter-pool]
- default-fc fc
- no default-fc
- description description-string
- no description
- dot1p-classification policy-id
- no dot1p-classification
- dscp-classification policy-id
- no dscp-classification
- [no] fc fc-name [create]
- broadcast-meter meter-id
- no broadcast-meter
- meter meter-id
- no meter
- multicast-meter meter-id
- no multicast-meter
- unknown-meter meter-id
- no unknown-meter
- [no] ip-mac-match {ip-first | mac-first}
- [no] ip-criteria [any | dscp-only]
- [no] entry entry-id [create]
- action [fc fc]
- no action
- description description-string
- no description
- match [protocol protocol-id]
- no match
- dscp dscp-value | dscp-name [dscp-mask]
- no dscp
- dst-ip {ip-address/mask | ip-address netmask}
```

```

- no dst-ip
- dst-port {eq} dst-port-number
- no dst-port
- ip-prec ip-prec-value [ip-prec-mask]
- no ip-prec
- src-ip {ip-address/mask | ip-address netmask}
- no src-ip
- src-port {eq} src-port-number
- no src-port
- renum [old-entry-id new-entry-id]
- [no] ipv6-criteria [any | dscp-only] [IPv6 Match Criteria]
- [no] entry entry-id [create]
- action [fc fc]
- no action
- description description-string
- no description
- match [next-header next-header]
- no match
- dscp dscp-value | dscp-name [dscp-mask]
- no dscp
- dst-ip {ipv6-address/prefix-length}
- no dst-ip
- dst-port {eq} dst-port-number}
- no dst-port
- ip-prec ip-prec-value [ip-prec-mask]
- no ip-prec
- src-ip {ipv6-address/prefix-length}
- no src-ip
- src-port {eq} src-port-number
- no src-port
- renum [old-entry-id new-entry-id]
- [no] mac-criteria [any | dot1p-only]
- [no] entry entry-id
- action [fc fc]
- no action
- description description-string
- no description
- [no] match
- dot1p dot1p-value [dot1p-mask]
- no dot1p
- dst-mac ieee-address [ieee-address-mask]
- no dst-mac
- etype 0x0600..0xffff
- no etype
- src-mac ieee-address [ieee-address-mask]
- no src-mac
- renum
- meter meter-id [multipoint] [create]
- no meter meter-id
- adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
- no adaptation-rule
- cbs size [kbits | bytes | kbytes]
- no cbs
- color-mode color-mode
- no color-mode
- mbs size [kbits | bytes | kbytes]
- no mbs
- mode {trtcm1 | trtcm2 | srtcm}
- no mode
- rate cir-rate-in-kbps [pir pir-rate-in-kbps]
- no rate
- num-qos-classifiers [num-resources] [ipv6 | no-ipv6]
- scope {exclusive | template}
- no scope

```

```
- table-classification-criteria table-classification-criteria
```

9.8.1.2 Table-based IP DSCP and dot1p classification policy commands for SAP ingress (7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)

```
config
- qos
- dot1p-classification classification-id [create]
- no dot1p-classification classification-id
- default-dot1p-fc fc profile [in | out | dei]
- no default-dot1p-fc
- description description-string
- no description
- dot1p dot1p-priority fc fc-name [profile {in | out | dei}]
- no dot1p dot1p-priority
- dscp-classification classification-id [create]
- no dscp-classification classification-id
- no default-dscp-fc fc profile {in | out}
- description description-string
- no description
- dscp dscp-name fc fc-name [profile {in | out}]
- no dscp dscp-name
```

9.8.1.3 Service meter commands for SAP ingress (7210 SAS-Mxp only)

```
config
- qos
- [no] sap-ingress policy-id [create] [use-svc-meter-pool]
- default-fc fc-name
- no default-fc
- dot1p-classification policy-id
- no dot1p-classification
- dscp-classification policy-id
- no dscp-classification
- fc-meter-map policy-id [create]
- no fc-meter-map policy-id
- counter-mode counter-mode
- no counter-mode
- fc fc-name [create]
- no fc fc-name
- broadcast-meter meter-id
- no broadcast-meter
- meter meter-id
- no meter
- multicast-meter meter-id
- no multicast-meter
- unknown-meter meter-id
- no unknown-meter
- meter meter-id [create]
- no meter meter-id
- adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
- no adaptation-rule
- cbs size [kbits | bytes | kbytes]
- no cbs
- color-mode color-mode
- no color-mode
- mbs size [kbits | bytes | kbytes]
```

```
- no mbs
- mode mode [trtcm1 | trtcm2 | srtcm]
- no mode
- rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
- [no] rate
```

9.8.1.4 Operational commands

```
config
- qos
- copy sap-ingress src-pol dst-pol [overwrite]
```

9.8.1.5 Show commands

```
show
- qos
- dot1p-classification [policy-id] association
- dot1p-classification [policy-id] [detail]
- dscp-classification [policy-id] association
- dscp-classification [policy-id] [detail]
- sap-ingress policy-id [detail | association | match-criteria]
- fc-meter-map [policy-id] association
- fc-meter-map [policy-id] [detail]
```

9.8.2 Command descriptions

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

9.8.2.1 Configuration commands

- [Generic commands](#)
- [Service ingress QoS policy commands](#)
- [Service ingress QoS policy forwarding class commands](#)
- [Service ingress QoS policy entry commands](#)
- [IP QoS policy match commands](#)
- [Service ingress MAC QoS policy match commands](#)
- [Service meter QoS policy commands](#)
- [IP DSCP and dot1p classification policy commands \(for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12\)](#)
- [Service meter commands for SAP ingress \(for 7210 SAS-Mxp\)](#)

9.8.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>dscp-classification

config>qos>dot1p-classification

config>qos>sap-ingress

config>qos>sap-ingress>ip-criteria>entry

config>qos>sap-ingress>ipv6-criteria>entry

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

Parameters

description-string

Specifies 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.

9.8.2.1.2 Service ingress QoS policy commands

sap-ingress

Syntax

[no] sap-ingress *policy-id* [create] [use-svc-meter-pool]

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates or edits the ingress policy. The ingress policy defines the Service Level Agreement (SLA) enforcement service packets receive as they ingress a SAP. SLA enforcement is accomplished through the definition of meters and queues (depends on the support available on a platform) that have Forwarding Class (FC), Committed Information Rate (CIR), Peak Information Rate (PIR), Committed Burst Size (CBS), and Maximum Burst Size (MBS) characteristics. The simplest policy defines a single queue or meter that all ingress traffic flows through. Complex policies have multiple meters/queues combined with classification entries that indicate which meter or queue a packet will flow through.

Policies in effect are templates that can be applied to multiple services as long as the **scope** of the policy is **template**. Meters defined in the policy are not instantiated until a policy is applied to a service SAP.

Only one service ingress policy can be provisioned. The SAP ingress policy with *policy-id* 1 is a system-defined policy applied to services when no other policy is explicitly specified. The system SAP ingress policy can be modified but not deleted. The **no sap-ingress** command restores the factory default settings when used on *policy-id* 1. See [Default SAP ingress policy](#) for more information.

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



Note:

- Before associating a SAP ingress policy with a SAP, resources must be allocated using the **config system resource-profile ingress-internal-tcam qos-sap-ingress-resource** command. For more information about this CLI command and resource allocation, see [Service ingress QoS policies](#) and the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide*.
- Only meters with the SAP ingress policy are supported on 7210 SAS-T (access-uplink and network mode), 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE; ingress queues are not supported.

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. The system default SAP ingress policy is a special case; the **no** command reverts the factory defaults to policy ID 1.



Note:

Default SAP ingress policy ID 1 is used in low-sap-scale mode. See the *7210 SAS-Mxp, S, Sx, T Services Guide* for more information about low-sap-scale mode and high-sap-scale mode.

Parameters

policy-id

Specifies the policy.

Values 1 to 65535

create

Keyword to create a SAP ingress policy.

use-svc-meter-pool

Specifies that the SAP ingress policy uses policers or meters from the service-meter pool. When this parameter is specified, an FC meter map is attached to the policy by default. Only table-based classification policies are used for classification; the CAM-based classification entries are ignored. If this parameter is not specified, the meter resources for the SAP ingress policy are allocated from the CAM-based meter resource pool. This parameter is only supported on 7210 SAS-Mxp.

scope

Syntax

scope {exclusive | template}

no scope

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the Service Ingress QoS policy scope as exclusive or template.

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

Default

template

Parameters

exclusive

Specifies that 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 becomes available for assignment to another exclusive SAP.

template

Specifies that the policy can be applied to multiple SAPs on the router.

Default QoS policies are configured with template scopes. An error is generated when the **template** scope parameter to **exclusive** scope on default policies is modified.

table-classification-criteria

Syntax

table-classification-criteria *table-classification-criteria*

Context

config>qos>sap-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command allows the user to choose the dot1p or DSCP classification policies to classify the traffic to an FC. The option to select either dot1p, DSCP, or both, applies only to SAPs configured in Layer 2 services (Epipe and VPLS). It is not applicable to SAPs configured in Layer 3 services (IES, VPRN and RVPLS). See [Table-based IP DSCP and dot1p classification policy commands for SAP ingress \(7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12\)](#) for more information.

The following options can be used to configure the classification policy.

- If **none** is configured, use **default-fc fc-name profile out** (from the SAP ingress policy).
- If **use-dscp** is configured, use the following policies:
 - DSCP classification policy for IP packets
 - **default-fc fc-name profile out** (from the SAP ingress policy) for non-IP packets
- If **use-dot1p** is configured, use the following policies:
 - dot1p classification policy for all tagged packets (IP and non-IP)
 - **default-fc fc-name profile out** (from the SAP ingress policy) for untagged packets
- If **both-dscp-dot1p** is configured, use the following policies:
 - DSCP classification policy for IP packets
 - dot1p classification policy for non-IP tagged packets
 - **default-fc fc-name profile out** (from the SAP ingress policy) for non-IP untagged traffic

Default

table-classification-criteria both-dscp-dot1p

Parameters

table-classification-criteria

Specifies the table classification criteria to use.

Values use-dscp, use-dot1p, both-dscp-dot1p, none

default-fc

Syntax

default-fc *fc*

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the default forwarding class for the policy. If an ingress packet does not match a higher priority (more explicit) classification command, the default forwarding class will be associated with the packet. Unless overridden by an explicit forwarding class classification rule, all packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class.

The default forwarding class is best effort (be). The **default-fc** settings are displayed in the **show configuration** and **save** output regardless of inclusion of the **detail** keyword.

Default

be

Parameters

fc

Specifies the forwarding class name for the queue or meter. The value given for *fc* must be one of the predefined forwarding classes in the system.

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

dscp-classification

Syntax

dscp-classification *policy-id*

no dscp-classification

Context

config>qos>sap-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command associates a DSCP classification policy with a SAP ingress QoS policy.

The **no** form of this command removes the DSCP classification policy from its association with the SAP ingress QoS policy.

Default

no dscp-classification

Parameters

policy-id

Specifies the policy ID that uniquely identifies the DSCP classification policy.

Values 1 to 65535

fc

Syntax

[no] fc *fc-name* [create]

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates a class instance of the forwarding class *fc-name*. Once the *fc-name* is created, classification actions can be applied and can be used in match classification criteria.

The **no** form of this command removes all the explicit queue mappings for *fc-name* forwarding types. The queue mappings revert to the default meters for *fc-name*.

Parameters

fc-name

Specifies the forwarding class name for the queue. The value given for the *fc-name* must be one of the predefined forwarding classes for the system.

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

create

Keyword to create a forwarding class.

ip-mac-match

Syntax

[no] ip-mac-match {ip-first | mac-first}

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command must be executed if user intends to match on both IP and MAC criteria in a SAP ingress policy. If this command is not executed software does not allow for configuration of both IP and MAC criteria in a SAP ingress policy. That is, without this command in a SAP ingress policy IP and MAC criteria are mutually exclusive.

The user also has the option to specify if all the IP criteria entries configured in the policy need to be matched first followed by all the MAC criteria entries or vice-versa. That is, if **ip-first** is configured all the IP criteria entries are matched and only if there are no matches in the MAC criteria entries are matched. If a match is found no further matches are done and the actions associated with the matched entry are taken.

Default

no ip-mac-match

Parameters

ip-first

Keyword to match all the IP criteria entries first before matching any of the MAC entries.

mac-first

Keyword to match all the MAC criteria entries first before matching any of the IP entries.

ip-criteria

Syntax

[no] **ip-criteria** *policy id* [**any** | **dscp-only**]

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates or edits policy entries that specify IP criteria DiffServ code point. IP criteria-based SAP ingress policies are used to select the appropriate ingress meter and corresponding forwarding class for matched traffic.

The user can specify either **any** or **dscp-only** as the sub-criteria. The sub-criteria determines what fields can be used to match traffic. The resource allocation for classification is affected by the sub-criteria in use. See [Table 62: SAP ingress resource allocation and match criteria types](#) for more information.

The 7210 SAS OS implementation exits on the first match found and executes the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.



Note:

When the **use-svc-meter-pool** parameter is set or when table-based-classification is used, the **ip-criteria** entries in the policy are ignored.

The **no** form of this command deletes all the entries specified under this node. Once IP criteria entries are removed from a SAP ingress policy, the IP criteria is removed from all services where that policy is applied.

Default

dscp-only

Parameters

any

Specifies that entries can use any of the fields available under ip-criteria for matching; for example, IP source, IP destination, IP protocol fields can be used.

dscp-only

Specifies that entries can use the IP DSCP field or IP precedence field.

policy-id

Specifies the policy.

Values 1 to 65535

ipv6-criteria

Syntax

[no] **ipv6-criteria** *policy-id* [**any** | **dscp-only**]

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command selects the appropriate ingress meters and corresponding forwarding class for matched traffic.

This command is used to enter the node to create or edit policy entries that specify IPv6 criteria such as IP quintuple lookup or DiffServ code point.

The 7210 SAS OS implementation exits on the first match found and executes the actions in accordance with the accompanying **action** command. For this reason entries must be sequenced correctly from most to least explicit.



Note:

Before associating a SAP ingress policy configured to use IPv6 criteria with a SAP, resources must be allocated using the **config system resource-profile ingress-internal-tcam qos-sap-ingress-resource ipv6-ipv4-match-enable** command. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information about this CLI command and resource allocation.

The **no** form of this command deletes all the entries specified under this node. Once ipv6-criteria entries are removed from a SAP ingress policy, the ipv6-criteria is removed from all services where that policy is applied.



Note:

When the **use-svc-meter-pool** parameter is set or when table-based-classification is used, the ip-criteria entries in the policy are ignored.

Parameters

any

Specifies that entries can use any of the fields available under **ipv6-criteria** for matching; for example, IPv6 source, IPv6 destination, IPv6 protocol fields can be used.

dscp-only

Specifies that entries can use the IP DSCP field or IPv6 precedence field.

policy-id

Specifies the policy.

Values 1 to 65535

mac-criteria

Syntax

[no] **mac-criteria** *policy id* [**any** | **dot1p-only**]

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command selects the appropriate ingress meters and corresponding forwarding class for matched traffic.

The user can specify either **any** or **dot1p-only** as the sub-criteria. The sub-criteria determines what fields can be used to match traffic. The resource allocation for classification is affected by the sub-criteria in use. See [Table 62: SAP ingress resource allocation and match criteria types](#) for more information.

This command is used to enter the node to create or edit policy entries that specify MAC criteria.

The 7210 SAS OS implementation exits on the first match found and executes the actions in accordance with the accompanying **action** command. For this reason, entries must be sequenced correctly from most to least explicit.



Note:

When the **use-svc-meter-pool** parameter is set or when table-based-classification is used, the ip-criteria entries in the policy are ignored.

The **no** form of this command deletes all the entries specified under this node. Once mac-criteria entries are removed from a SAP ingress policy, the mac-criteria is removed from all services where that policy is applied.

Default

any

Parameters

any

Specifies that entries can use any of the fields available under mac-criteria; for example, MAC source, MAC destination, and MAC Ethertype fields can be used)

dot1p-only

Specifies that entries can use only the dot1p field.

policy-id

Specifies the policy.

Values 1 to 65535

num-qos-classifiers

Syntax

num-qos-classifiers [*num-resources*] [**ipv6** | **no-ipv6**]

no num-qos-classifiers

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the number of classifiers the SAP ingress Qos policy can use. A user cannot modify this parameter when it is in use (that is, applied to a SAP).

The *num-resources* parameter determines the maximum number of meters that are available to this policy. The maximum number of meters available for use by the forwarding classes (FC) defined under this policy is equal to half the value specified in the parameter *num-resources*. Any of these meters is available for use to police unicast or multipoint traffic. Any of these meters is available for use by more than one FC (or a single meter is available for use by all the FCs).

The keyword **ipv6** allows users to indicate that they plan to use the ipv6-criteria and the resources needed for this SAP ingress QoS policy must be allocated for the chunk allocated to IPv6 criteria.

On the 7210 SAS-Mxp, this parameter is ignored if the **use-svc-meter-pool** parameter is set.

Default

num-qos-classifiers 2 no-ipv6

Parameters

num-resources

Specifies the number of resources planned for use by this policy, expressed as multiple of 2.

Values 2 to 256

Default 2

ipv6

Specifies that the user intends to use the ipv6-criteria and software must allocate resources from the chunks allotted to IPv6 criteria.

no-ipv6

Specifies that the user intends to use the ipv6-criteria. Resources are then allocated from the chunk allotted to either IPv4 criteria or MAC criteria, depending on what criteria the user uses.

9.8.2.1.3 Service ingress QoS policy entry commands

action

Syntax

action [*fc fc*]

no action

Context

config>qos>sap-ingress>ip-criteria>entry

config>qos>sap-ingress>ipv6-criteria>entry

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This mandatory command associates the forwarding class with specific IP or MAC criteria entry ID. The **action** command supports setting the forwarding class parameter. Packets that meet all match criteria within the entry have their forwarding class overridden based on the parameters included in the **action** parameters.

The **action** command must be executed for the match criteria to be added to the active list of entries.

Each time action is executed on a specific entry ID, the previous entered values for **fc** is overridden with the newly defined parameters.

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.

Default

action specified by the **default-fc**.

Parameters

fc *fc*

Specifies the forwarding class name for the queue or meter. The value given for *fc* must be one of the predefined forwarding classes in the system.

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

entry

Syntax

[no] **entry** *entry-id* [create]

Context

config>qos>sap-ingress>ip-criteria

config>qos>sap-ingress>ipv6-criteria

config>qos>sap-ingress>mac-criteria

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates or edits an IP or MAC criteria entry for the policy. Multiple entries can be created using unique *entry-id* numbers.

The list of flow criteria is evaluated in a top down fashion with the lowest entry ID at the top and the highest entry ID at the bottom. If the defined match criteria for an entry within the list matches the information in the egress packet, the system stops matching the packet against the list and performs the matching entries reclassification actions. If none of the entries match the packet, the IP flow reclassification list has no effect on the packet.

An entry is not populated in the list unless the **action** command is executed for the entry. An entry that is not populated in the list has no effect on egress packets. If the **action** command is executed without any explicit reclassification actions specified, the entry is populated in the list allowing packets matching the entry to exit the list, preventing them from matching entries lower in the list. Because this is the only flow reclassification entry that the packet matched and this entry explicitly states that no reclassification action is to be performed, the matching packet will not be reclassified.

The **no** form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services where that policy is applied.

Parameters

entry-id

Specifies a match criterion and the corresponding action. It is recommended that multiple entries be given *entry-ids* in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword **action fc fc-name** for it to be considered complete. Entries without the **action** keyword will be considered incomplete and therefore will be rendered inactive.

Values 1 to 64

create

Keyword to create a flow entry when the system is configured. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the flow entry already exists.

match

Syntax

[no] match [**protocol** *protocol-id*]

Context

config>qos>sap-ingress>ip-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures match criteria for SAP QoS policies. When the match criteria have been satisfied the action associated with the match criteria is executed.

Only a single match criteria (either MAC or IP) is allowed at any point of time.

Parameters

protocol *protocol-id*

Specifies an IP protocol to be used as a SAP QoS policy match criterion.

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

Values 0 to 255

match

Syntax

match

no match

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables the context for entering and editing match MAC criteria for ingress SAP QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, all criteria must be satisfied (AND function) before the action associated with the match will be executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

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

match

Syntax

match [*next-header next-header*]

no match

Context

config>qos>sap-ingress>ipv6-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures match criteria for ingress SAP QoS policy match IPv6 criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, then all criteria must be satisfied (AND function) before the action associated with the match is executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

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

Parameters

next-header *next-header*

Specifies the next header to match.

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

Values protocol numbers accepted in decimal, hexadecimal, or binary: 0 to 42, 45 to 49, 52 to 59, 61 to 255

keywords — none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, cmp, idrp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, iso-p, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

* — udp/tcp wildcard

9.8.2.1.4 Service ingress QoS policy forwarding class commands

broadcast-meter

Syntax

broadcast-meter *meter-id*

no broadcast-meter

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

The broadcast forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of this command reverts the broadcast forwarding type *meter-id* to the default of tracking the multicast forwarding type meter mapping.

Parameters

meter-id

Specifies an existing multipoint queue defined in the **config>qos>sap-ingress** context.

Values 2 to 32

Default 11

meter

Syntax

meter *meter-id*

no meter

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command overrides the default unicast forwarding type meter mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy as a non-multipoint meter before the mapping can be made. Once the forwarding class mapping is executed, all unicast traffic (this includes all traffic, even broadcast and multicast for services) on a SAP using this policy is forwarded using the *meter-id*.

The **no** form of this command reverts the unicast (point-to-point) *meter-id* to the default meter for the forwarding class (meter 1).

Parameters

meter-id

Specifies an existing non-multipoint meter defined in the **config>qos>sap-ingress** context.

Values 1 to 32

multicast-meter

Syntax

multicast-meter *meter-id*

no multicast-meter

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

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

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

Parameters

meter-id

Specifies an existing multipoint queue defined in the **config>qos>sap-ingress** context.

Values 1 to 32

Default 11

unknown-meter

Syntax

unknown-meter *meter-id*

no unknown-meter

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command overrides the default unknown unicast forwarding type meter mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy as a multipoint meter before the mapping can be made. After the forwarding class mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the *meter-id*.

The unknown forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of this command reverts the unknown forwarding type *meter-id* to the default of tracking the multicast forwarding type meter mapping.

Parameters

meter-id

Specifies an existing multipoint meter defined in the **config>qos>sap-ingress** context.

Values 1 to 32

Default 11

9.8.2.1.5 IP QoS policy match commands

dscp

Syntax

dscp *dscp-value* | *dscp-name* [*dscp-mask*]

no dscp

Context

config>qos>sap-ingress>ip-criteria>entry>match

config>qos>sap-ingress>ipv6-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures a DiffServ Code Point (DSCP) code point to be used for classification of packets from the specified FC.

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



Note:

- This feature is applicable for ip-criteria (any and dscp-only), and ipv6-criteria (any and dscp-only).
- Either DSCP name or DSCP value with a mask can be configured.
- When the user configures *dscp* value alone, the **show** command displays the *dscp* value as the configured value and the *dscp-mask* as the FC.

Parameters

dscp-value

Specifies the DSCP value in either hexadecimal format, decimal format or binary format.

Values 0 to 64

dscp-name

Specifies a DSCP name that has been previously mapped to a value using the **dscp-name** command. The DiffServ code point can only be specified by its name.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

dscp-mask

Specifies a 6-bit mask that 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 4 and 0b000100 for value and mask.

Values 0 to 64 (decimal, hexadecimal, or binary)

Default 64 (exact match)

dst-ip

Syntax

dst-ip {*ip-address/mask* | *ip-address netmask*}

no dst-ip

Context

config>qos>sap-ingress>ip-criteria>entry>match
config>qos>sap-ingress>ipv6-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures a destination address range to be used as a SAP QoS policy match criterion.

To match on the destination address, specify the address and its associated mask; for example, 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

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

Parameters

ip-address

Specifies the IP address of the destination IP or IPv6 interface. This address must be unique within the subnet and specified in dotted-decimal notation.

Values ipv4-prefix — a.b.c.d
 ipv6-prefix — x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:d.d.d.d
 x — 0 to FFFF (hexadecimal)
 d — 0 to 255 (decimal)
 ipv4-prefix-length — 0 to 32
 ipv6-prefix-length — 0 to 128

netmask

Specifies the subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255

dst-port

Syntax

dst-port {eq} *dst-port-number*

no dst-port

Context

config>qos>sap-ingress

config>qos>sap-ingress>ip-criteria>entry>match

config>qos>sap-ingress>ipv6-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures a destination TCP or UDP port number for a SAP QoS policy match criterion.

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

Parameters

eq dst-port-number

Specifies the TCP or UDP port number to match, specified as equal to (**eq**) the destination port value specified as a decimal integer.

Values 1 to 65535 (decimal hex or binary)

fragment

Syntax

fragment {true | false}

no fragment

Context

config>qos>sap-ingress>ip-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures fragmented or non-fragmented IP packets as a SAP QoS policy match criterion. The **no** form of this command removes the match criterion.

Default

fragment false

Parameters

true

Keyword to configure 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

Keyword to configure 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.

ip-prec

Syntax

ip-prec *ip-prec-value* [*ip-prec-mask*]

no ip-prec

Context

config>qos>sap-ingress>ip-criteria>entry>match

config>qos>sap-ingress>ipv6-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command defines a specific IP Precedence value that must be matched to perform the associated classification actions. If an ingress packet on the SAP where the SAP ingress QoS policy is applied to matches the specified IP Precedence value, the actions associated with this entry are taken.

The *ip-prec-value* is derived from the most significant three bits in the IP header ToS byte field (precedence bits). The three precedence bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop QoS behavior. The precedence bits are also part of the newer DiffServ Code Point (DSCP) method of mapping packets to QoS behavior. The DSCP uses the most significant six bits in the IP header ToS byte and so overlaps with the precedence bits.

Both IP precedence and DSCP classification rules are supported. A match entry cannot match on both IP DSCP and IP precedence values. That is, the user can use either IP DSCP or IP precedence match in a match entry but not both. The software blocks configuration of IP precedence match if **ip-dscp** is configured already. The converse is also true. A single policy having multiple match entries can have entries such that some of them match IP DSCP and some others match IP precedence. The order of the entry determines the priority of the match.

The **no** form of this command removes the IP Precedence match criterion.

Parameters

ip-prec-value

Specifies the unique IP header ToS byte precedence bits value that will match the IP precedence rule.

Values 0 to 7

ip-prec-mask

Specifies the mask to use for the match.

src-ip

Syntax

src-ip {*ip-address/mask* | *ip-address netmask*}

no src-ip

Context

config>qos>sap-ingress>ip-criteria>entry>match

config>qos>sap-ingress>ipv6-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures a source IP or IPv6 address range to be used as an SAP QoS policy match criterion.

To match on the source IP or IPv6 address, specify the address and its associated mask; for example, 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

The **no** form of this command removes the source IP or IPv6 address match criterion.

Parameters

ip-address | ipv6-address

Specifies the IP or IPv6 address of the source IP interface. This address must be unique within the subnet and specified in dotted-decimal notation.

Values ipv4-prefix — a.b.c.d
 ipv6-prefix — x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:d.d.d.d
 x — 0 to FFFF (hexadecimal)
 d — 0 to 255 (decimal)
 ipv4-prefix-length — 1 to 32
 ipv6-prefix-length — 0 to 128

mask

Specifies the subnet mask length, expressed as an integer or in dotted-decimal notation.

Values 0 to 32

netmask

Specifies the subnet mask, in dotted-decimal notation.

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

src-port

Syntax

src-port {eq} *src-port-number*

no src-port

Context

config>qos>sap-ingress>ip-criteria>entry>match

config>qos>sap-ingress>ipv6-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures a source TCP or UDP port number for a SAP QoS policy match criterion.
The **no** form of this command removes the source port match criterion.

Parameters

eq src-port-number
Specifies the TCP or UDP port number to match specified as equal to (**eq**) to the source port value specified as a decimal integer.

Values 1 to 65535 (decimal, hexadecimal, or binary)

9.8.2.1.6 Service ingress MAC QoS policy match commands

dot1p

Syntax

dot1p dot1p-value [dot1p-mask]
no dot1p

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the IEEE 802.1p value to be used as the match criterion.
The **no** form of this command removes the dot1p value as the match criterion.

Parameters

dot1p-value
Specifies the IEEE 802.1p value in decimal.

Values 0 to 7

dot1pmask
Specifies a 3-bit mask that can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	D	4
Hexadecimal	0xH	0x4

Format Style	Format Syntax	Example
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 (decimal) (exact match)

dst-mac

Syntax

dst-mac *ieee-address* [*ieee-address-mask*]
no dst-mac

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures a destination MAC address or range to be used as a Service Ingress QoS policy match criterion.

The **no** form of this command removes the destination mac address as the match criterion.

Parameters

ieee-address

Specifies the MAC address to be used as a match criterion.

Values	HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit
--------	---

ieee-address-mask

Specifies a 48-bit mask to match a range of MAC address values.

This 48-bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHHHH	0x0FFFFFF000000
Binary	0BBBBBBB...B	0b11110000...B

All packets with a source MAC OUI value of 00-03-FA subject to a match condition should be specified as: 0003FA000000 0x0FFFFFF000000

Values 0x0000000000000000 to 0xFFFFFFFFFFFFFFF (hex)

Default 0xFFFFFFFFFFFFFFF (hex) (exact match)

etype

Syntax

etype *etype-value*

no etype

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures an Ethernet type II value for use as a service ingress QoS policy match criterion.

The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame. For example, 0800 is used to identify IPv4 packets.

The Ethernet type field is used by the Ethernet version-II frames. IEEE 802.3 Ethernet frames use the dsap, ssap or snap-pid fields as match criteria; the Ethernet type field is not used.

The snap-pid, etype, ssap, and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The dataplane processes a maximum of two VLAN tags in a received packet. The Ethertype used in the MAC matching criteria for ACLs is the Ethertype that is found in the packet after processing single-tagged frames, double-tagged frames, and no-tag frames

The packet is considered to have no tags if at least one of the following criteria is true:

- the packet is a null-tagged frame
- the packet is a priority-tagged frame
- the outermost Ethertype does not match the default Ethertype (0x8100)
- the outermost Ethertype does not match the configured dot1q-etype on Dot1q encapsulated ports
- the outermost Ethertype does not match the configured qinq-etype on QinQ encapsulated ports

The packet is considered to have a single tag if at least one of the following criteria is true:

- the outermost Ethertype matches the default Ethertype (0x8100)
- the outermost Ethertype matches the configured dot1q-etype on Dot1q encapsulated ports
- the outermost Ethertype matches the configured qinq-etype on QinQ encapsulated ports

The packet is considered to have double tags if at least one of the following criteria is true:

- the outermost Ethertype matches the default Ethernet type (0x8100)
- the configured dot1q-etype on Dot1q encapsulated ports and the immediately following Ethertype match the default Ethertype (0x8100)
- the configured qinq-etype on QinQ encapsulated ports and the immediately following Ethertype match the default Ethertype (0x8100)

The **no** form of this command removes the previously entered etype field as the match criteria.

Parameters

etype-value

Specifies the Ethernet II frame Ethertype value to be used as a match criterion, expressed in decimal or hexadecimal.

Values 0x0600 to 0xFFFF (1536 to 65535)

src-mac

Syntax

src-mac *ieee-address* [*ieee-address-mask*]

no src-mac

Context

config>qos>sap-ingress>mac-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures a source MAC address or range to be used as a service ingress QoS policy match criterion.

The **no** form of this command removes the source mac as the match criteria.

Parameters

ieee-address

Specifies the 48-bit IEEE MAC address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask

Specifies a 48-bit mask. This 48 bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHH	0x0FFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

To configure all packets with a source MAC OUI value of 00-03-FA are subject to a match condition, then the entry should be specified as: 003FA000000 0xFFFFFFFF000000

Values 0x0000000000000000 to 0xFFFFFFFFFFFFFFF (hexadecimal)

Default 0xFFFFFFFFFFFFFFF (hexadecimal) (exact match)

9.8.2.1.7 Service meter QoS policy commands

meter

Syntax

meter *meter-id* [**multipoint**] [**create**]

no meter *meter-id*

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure an ingress service access point (SAP) QoS policy meter.

This command allows the creation of multipoint meters. Only multipoint meters 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 meters special handling of the multipoint traffic is possible. Each meter acts as an accounting and (optionally) policing 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 meter based on forwarding type) needs to be defined. The individual classification rules used to place traffic into forwarding classes are not affected. Meters must be defined as multipoint at the time of creation within the policy.

The multipoint meters 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 meter.

When an ingress SAP QoS policy with multipoint meters is applied to an Epipe SAP, the multipoint meters are not created.

Any billing or statistical queries about a multipoint meter on a non-multipoint service returns zero values. Any meter parameter information requested about a multipoint meter on a non-multipoint service returns the meter parameters in the policy. Multipoint meters would not be created for non-multipoint services.

The **no** form of this command removes the *meter-id* from the SAP ingress QoS policy and from any existing SAPs using the policy. Any forwarding class mapped to the meter, will revert to their default meters. When a meter is removed, any pending accounting information for each SAP meter created due to the definition of the meter in the policy is discarded.

Parameters

meter-id

Specifies the meter ID, expressed as an integer. The *meter-id* uniquely identifies the meter within the policy. This is a required parameter each time the meter command is executed.

Values 1 to 32

multipoint

Keyword that creates the meter as a multipoint meter.

create

Keyword to create a meter.

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]

no adaptation-rule

Context

config>qos>sap-ingress>meter

config>qos>fc-meter-map>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>qos>fc-meter-map>meter** context is only supported on the 7210 SAS-Mxp.

This command defines the method used by the system to derive the operational CIR and PIR rates when the meter is provisioned in hardware. For the **cir** and **pir** parameters, the system attempts to find the best operational rate depending on the defined constraint.



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

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. When a specific **adaptation-rule** is removed, the default constraints for **cir** and **pir** apply.

Default

adaptation-rule cir closest pir closest

Parameters

cir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced to adapt the CIR rate defined using the **meter meter-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used to derive the operational CIR rate for the meter. When the **cir** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

Values **max** — Specifies that the operational CIR value is equal to or less than the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational CIR value is equal to or greater than the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

pir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced to adapt the PIR rate defined using the **meter meter-id rate** command. The **pir** parameter requires a qualifier that defines the constraint used to derive the operational PIR rate for the meter. When the **rate** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

Values **max** — Specifies that the operational PIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational PIR value is equal to or greater than the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

cbs

Syntax

cbs *size* [kbits | bytes | kbytes]

no cbs

Context

config>qos>sap-ingress>meter

config>qos>fc-meter-map>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>qos>fc-meter-map>meter** context is only supported on the 7210 SAS-Mxp.

This command provides a mechanism to override the default CBS for the meter. The committed burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the CBS value then the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

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

Default

cbs 32 kbits

Parameters

size

Specifies the number of kilobits, kilobytes, or bytes that are reserved for the meter. For example, if a value of 100 kbits is required, enter the value 100. The bucket size is rounded off to the next highest 4096 bytes boundary.

Values **kbits** — 4 to 2146959, default
 bytes — 512 to 274810752, default
 kbytes — 1 to 268369, default

color-mode

Syntax

color-mode *color-mode*

no color-mode

Context

config>qos>sap-ingress>meter

config>qos>fc-meter-map>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>qos>fc-meter-map>meter** context is only supported on the 7210 SAS-Mxp.

This command specifies that the meter operates in either color-aware mode or color-blind mode.

In color-blind mode, the profile/color assigned to the packet on ingress is ignored. The CIR and PIR rate configured for the meter is used to determine the final color/profile for the packet. If the packet is within the CIR, then the final profile/color assigned to the packet is in-profile/green and if the packets exceeds the CIR and is within the PIR, then the final profile/color assigned to the packet is out-of-profile/yellow. Packets that exceed the PIR rate are dropped.

In color-aware mode, the meter uses the profile assigned to the packet on ingress. Assign the profile on ingress by doing one of the following:

- enable DEI classification on access ports
- enable DSCP classification on access ports (7210 SAS-Mxp only)
- assign profile based on either Dot1p or DEI as done on network ports and access-uplink ports

In color-aware mode, the following behavior is expected.

- If the packet is pre-colored as an "in-profile" packet (which are also called "green" packets), depending on the burst size of the packet, the meter can mark the packet in-profile or out-profile.
- If the packet is pre-colored as an "out-profile" packet (which are also called "yellow" packets), even if the packet burst is less than the current available CBS, it would not be marked as in-profile and remain as out-profile.
- If the packet burst is higher than the MBS then it would be marked as "red" and would be dropped by meter at ingress.



Note:

The final disposition of the packet when a hierarchical meter is used depends on the parent meter/policer that is associated with the per FC meter. For example, on the 7210 SAS, when a SAP ingress aggregate meter is used, the final color, and therefore the action taken, also depends on the rate that is configured for the SAP ingress aggregate meter. The SAP ingress aggregate meter, which acts as the parent meter, accounts for the color assigned by the per FC meter.

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

Default

color-blind

Parameters

color-mode

Specifies if the meter operates in color-aware or color-blind mode.

Values color-blind, color-aware

mbs

Syntax

mbs *size* [kbits | bytes | kbytes]

no mbs

Context

config>qos>sap-ingress>meter

config>qos>fc-meter-map>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>qos>fc-meter-map>meter** context is only supported on the 7210 SAS-Mxp.

This command provides a mechanism to override the default MBS for the meter. The **mbs** command specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the MBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter-configured parameters.



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

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

Default

mbs 512 kbits

Parameters

size

Specifies the number of kilobits, kilobytes, or bytes that are reserved for the meter. For example, if a value of 100 kbits is required, enter the value 100. The bucket size is rounded off to the next highest 4096 bytes boundary.

Values	kbits — 4 to 2146959
	bytes — 512 to 274810752
	kbytes — 1 to 268369

mode

Syntax

mode {trtcm1 | trtcm2 | srtcm}

no mode

Context

config>qos>sap-ingress>meter

config>qos>fc-meter-map>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>qos>fc-meter-map>meter** context is only supported on the 7210 SAS-Mxp.

This command defines the mode of the meter. The mode can be configured as Two Rate Three Color Marker (trTCM1) or Single Rate Three Color Marker (srTCM). The mode command can be executed at any time.



Note:

- The meter counters are reset to zero when the meter mode is changed.
- For more information about the interpretation of rate parameters when the meter mode is configured as **trtcm2**, see the command description of the policer [rate](#) command.

The **no** form of this command reverts to the default mode **trtcm1**.

Default

mode trtcm 1

Parameters

trtcm1

Keyword to implement the policing algorithm defined in RFC2698. This keyword meters the packet stream and marks its packets either green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked either yellow or green depending on whether it exceeds the CIR. The trTCM1 is useful, for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate. Two token buckets are used, the CBS bucket and the MBS bucket. Tokens are added to the buckets based on the CIR and PIR rates. The algorithm deducts tokens from both the CBS and the MBS buckets to determine a profile for the packet.

trtcm2

Keyword to implement the policing algorithm defined in RFC4115. This keyword meters the packet stream and marks its packets either green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked either yellow or green depending on whether it exceeds or does not exceed the CIR. The trtcm2 is useful; for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate. Two token buckets are used, the CBS bucket and the EBS bucket. Tokens are added to the buckets based on the CIR and EIR rates. The algorithm deducts tokens from either the CBS bucket (that is, when the algorithm identifies the packet as in-profile or green packet) or the EBS bucket (that is, when the algorithm identifies the packet as out-of-profile or yellow packet).



Note:

When the meter mode is configured in **trtcm2** mode, the system interprets the PIR rate parameter as EIR for use by the RFC 4115 algorithm.

srtcm

Keyword to meter an IP packet stream and marks its packets either green, yellow, or red. Marking is based on a CIR and two associated burst sizes, a CBS and an Maximum Burst Size (MBS). A packet is marked green if it does not exceed the CBS, yellow if it does exceed the CBS, but not the MBS, and red otherwise. The srTCM is useful, for example, for ingress policing of a service, where only the length, not the peak rate, of the burst determines service eligibility.

rate

Syntax

rate *cir* *cir-rate-in-kbps* [**pir** *pir-rate-in-kbps*]

no rate

Context

config>qos>sap-ingress>meter

config>qos>fc-meter-map>meter

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>qos>fc-meter-map>meter** context is only supported on the 7210 SAS-Mxp.

This command defines the administrative PIR and CIR parameters for the meter.

It alters the PIR and CIR for all meters created through the association of the SAP ingress QoS policy with the *meter-id*.



Note:

If the **trtcm2** keyword is used to configure the **meter mode** commands, the system interprets the PIR parameter as the EIR for use by the algorithm defined in RFC 4115. Consequently, the system configures the policer EIR based on the configured PIR value.

The **no** form of this command reverts all meters created with the specific *meter-id* by through the association of the QoS policy with the default PIR and CIR parameters (**max**, 0). The **max** default specifies the amount of bandwidth in kilobits per second.

The following table lists the maximum CIR and PIR values for 7210 SAS platforms.

Table 63: Maximum CIR and PIR values for 7210 SAS platforms

7210 SAS platform	Maximum meter rate value (in Gb/s)
7210 SAS-Mxp	40 Gb/s
7210 SAS-R6	100 Gb/s
7210 SAS-R12	100 Gb/s
7210 SAS-Sx/S 1/10GE	40 Gb/s
7210 SAS-Sx 10/100GE	400 Gb/s
7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE in standalone-VC mode	40 Gb/s
7210 SAS-T	20 Gb/s

Default

rate cir 0 pir max

Parameters

cir *cir-rate-in-kbps*

Specifies the CIR, which overrides the default administrative CIR used by the meter. If the **rate** command has not been executed or the *cir-rate-in-kbps* parameter value is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and the CIR must be specified as a positive integer.

The actual CIR depends on the **meter adaptation-rule** command parameters and the hardware.

Values 0 to 20000000, **max**

pir *pir-rate-in-kbps*

Specifies the administrative PIR, in kilobits, for the meter. A valid PIR setting must be explicitly defined when this command is run. If the **rate** command has not been executed, the default PIR of **max** is assumed. If the **rate** command is executed, setting the PIR is optional. The **max** value is mutually exclusive with the *pir-rate-in-kbps* value.

Fractional values are not allowed and the PIR must be specified as a positive integer.

The actual PIR depends on the **meter adaptation-rule** command parameters and the hardware.

Values 0 to 20000000, **max**

9.8.2.1.8 IP DSCP and dot1p classification policy commands (for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12)

dot1p-classification

Syntax

dot1p-classification *classification-id* [**create**]

no dot1p-classification *classification-id*

Context

config>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command defines the map from the dot1p value that is in the Ethernet header of the received frame to the FC and ingress profile.

The **no** form of this command removes the definition of the map. The **no** form of this command replaces the specified dot1p-classification with dot1p-classification 1 in the SAP ingress policy or access port.

Default

default policy 1

Parameters

classification-id

Specifies the unique identifier for the policy.

Values 1 to 65535

create

Keyword to create the dot1p classification.

default-dot1p-fc

Syntax

default-dot1p-fc *fc* **profile** {*in* | *out* | *dei*}

no default-dot1p-fc

Context

config>qos>dot1p-classification

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command assigns the configured FC and profile to all VLAN tagged Ethernet packets that do not match any of the explicitly configured dot1p values.

The **no** form of this command assigns a default FC value of **be** and a profile value of **out**.

Default

default-dot1p-fc be profile out

Parameters

fc

Specifies the forwarding class.

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

profile {*in* | *out* | *dei*}

Specifies that all packets assigned to this forwarding class are considered in or out of profile based on this command. A value must be specified when the **profile** keyword is used in the command. If the profile is not assigned to a forwarding class, the packets of that FC are treated as "out-of-profile" packets.

Values **in** — Defines the packet profile as "in-profile"

out — Defines the packet profile as "out-of-profile"

dei — Specifies that DEI is used to determine the initial profile of the packet

dot1p

Syntax

```
dot1p dot1p-priority fc fc-name [profile {in | out | dei}]  
no dot1p dot1p-priority
```

Context

config>qos>dot1p-classification

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command maps the dot1p value to a forwarding class and also assign the ingress profile to the packet. This command can be repeated for each dot1p value that the user needs to map. The configured value is used to match the value in the outermost VLAN tag of the received tagged Ethernet packet and assign the configured forwarding class and profile on a exact match.

The **no** form of this command removes the mapping of the dot1p value to forwarding class.

A default forwarding class is not assigned on executing the **no** form of this command. If the VLAN tagged packet does not match the explicitly configured dot1p values, it is assigned the FC and profile value configured with the **default-dot1p-fc** command.

Default

default forwarding class is not assigned and must be explicitly configured

Parameters

fc

Specifies the forwarding class.

fc-name

Specifies the system-defined forwarding class name (it is case-sensitive).

profile {in | out | dei}

Specifies that all packets assigned to this forwarding class are considered in or out of profile based on this command. A value must be specified when the **profile** keyword is used in the command. If the profile is not assigned to a forwarding class, the packets of that FC are treated as "out-of-profile" packets.

Values

in — Defines the packet profile as "in-profile"

out — Defines the packet profile as "out-of-profile"

dei — Specifies that DEI is used to determine the initial profile of the packet

dot1p-priority

Specifies the dot1p priority value to match.

Values 0 to 7

dscp-classification

Syntax

dscp-classification *classification-id* [create]

no dscp-classification *classification-id*

Context

config>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command creates or edits a DSCP classification policy.

The **no** form of this command replaces the specified DSCP classification policy with dscp-classification 1 in the respective SAP ingress policy or access port.

Default

default policy 1

Parameters

classification-id

Specifies the ID of the DSCP classification policy.

Values 1 to 65535

default-dscp-fc

Syntax

default-dscp-fc *fc* profile {in | out}

[no] **default-dscp-fc**

Context

config>qos>dscp-classification

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures the default FC and profile values used to map all DSCP values that are not explicitly defined using the **dscp** command.

The **no** form of this command reverts the configured **default-dscp-fc** command FC and profile values to their default settings.

Default

default-dscp-fc "be" profile out

Parameters

fc

Specifies the forwarding class assigned as the default FC.

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

Default be

profile {in | out}

Specifies the profile state assigned as the default profile, either in-profile or out-of-profile.

Default out

dscp

Syntax

dscp *dscp-name* **fc** *fc-name* [**profile** {in | out}]

no dscp *dscp-name*

Context

config>qos>dscp-classification

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command maps a DSCP value to an FC and profile. The command can be repeated for each DSCP value that the user needs to map. It is used to match the IP DSCP in the received IP packets to assign the configured FC and profile value.

The **no** form of this command removes the mapping for the specified *dscp-name*. If the IP packet does not match the explicitly configured IP DSCP values, it is assigned the FC and profile value that is configured with the **default-dscp-fc** command.

Default

no dscp

Parameters

dscp-name

Specifies the IP DSCP value that gets mapped to the FC and profile specified in the command.

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-name

Specifies the name of the forwarding class assigned to the *dscp-name*.

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

in | out

Specifies the profile state assigned to the *dscp-name*, either in-profile or out-of-profile.

9.8.2.1.9 Service meter commands for SAP ingress (for 7210 SAS-Mxp)

dot1p-classification

Syntax

dot1p-classification *policy-id*

no dot1p-classification

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies that the SAP ingress policy uses the dot1p classification for the specified policy ID.

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

Default

policy-id 1

Parameters

policy-id

Specifies the SAP ingress policy ID that will use the dot1p classification.

Values 1 to 65535

dscp-classification

Syntax

dscp-classification *policy-id*

no dscp-classification

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies that the SAP ingress policy uses the DSCP classification for the specified policy ID.

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

Default

policy id 1

Parameters

policy-id

Specifies the DSCP classification policy ID that the SAP ingress policy will use.

Values 1 to 65535

fc-meter-map

Syntax

fc-meter-map *policy-id* [**create**]

no fc-meter-map *policy-id*

Context

config>qos

Platforms

7210 SAS-Mxp

Description

This command specifies an FC meter map for a SAP ingress policy. The user cannot modify the FC meter map if it is being referenced by a SAP. The user must update the map before attaching SAP ingress policies on the SAPs.



Note:

A single user configurable **fc-meter-map** policy is available to be used by all SAP ingress policies.

The **no** form of this command reverts to the default FC meter map policy.

Default

fc-meter-map 1

Parameters

policy-id

Specifies the policy ID for the FC meter map policy.

Values 1

create

Specifies an FC meter map policy that is not the default.

counter-mode

Syntax

counter-mode *counter-mode*

no counter-mode

Context

config>qos>fc-meter-map

Platforms

7210 SAS-Mxp

Description

This command is used to configure a counter mode for an FC meter map policy. The **in-out-profile-count** parameter specifies that in-profile and out-profile packets and octets are counted per meter. The **forward-drop-count** parameter specifies that forwarded and dropped packets and octets are counted per meter. For all counter modes, packets and bytes are counted simultaneously. The counter counts the in-profile and out-profile packets and octets based on the final profile assigned by the meter/policer using the configured rates.

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

Default

in-out-profile-count

Parameters

counter-mode

Specifies which counter mode the FC meter map policy uses.

Values in-out-profile-count, forward-drop-count

fc

Syntax

fc *fc-name* [**create**]

no fc *fc-name*

Context

config>qos>fc-meter-map

Platforms

7210 SAS-Mxp

Description

This command configures FC to meter mappings for an FC meter map policy.

The **no** form of this command removes all the explicit FC to meter mappings and reverts to the default.

Default

fc be

Parameters

fc-name

Specifies the forwarding class that is assigned as the default FC.

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

create

Keyword to create an FC for an FC meter map policy that is not the default.

broadcast-meter

Syntax

broadcast-meter *meter-id*

no broadcast-meter

Context

config>qos>fc-meter-map>fc

Platforms

7210 SAS-Mxp

Description

This command overrides the default broadcast forwarding type meter mapping for the **fc** *fc-name* command.

The *meter-id* parameter must exist within the policy before the mapping can be made. After the forwarding class mapping is executed, all broadcast traffic on a SAP that is using this policy is forwarded using the *meter-id* parameter specified.

The broadcast forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior. For more information, see [Use of the keyword "multipoint" for default meter "11"](#) and [Service ingress meter selection rules](#).

The **no** form of this command reverts the broadcast forwarding type *meter-id* parameter to the default of tracking the multicast forwarding type meter mapping.

Parameters

meter-id

Specifies the meter ID for the broadcast meter.

Values 1 to 32

Default 1

meter

Syntax

meter *meter-id*

no meter

Context

config>qos>fc-meter-map>fc

Platforms

7210 SAS-Mxp

Description

This command overrides the default unicast forwarding type meter mapping for the **fc** *fc-name* command. The specified *meter-id* parameter must exist within the policy before the mapping can be made. After the

forwarding class mapping is executed, all unicast traffic on a SAP using this policy is forwarded using the *meter-id* parameter specified.

If meter 11 is not configured in the policy, all multipoint traffic (that is, multicast, broadcast, and unknown-unicast traffic) is forwarded using the *meter-id*. For more information, see [Use of the keyword "multipoint" for default meter "11"](#) and [Service ingress meter selection rules](#).

The **no** form of this command reverts the unicast *meter-id* to the default meter for the FC.

Parameters

meter-id

Specifies the meter ID for the meter.

Values 1 to 32

Default 1

multicast-meter

Syntax

multicast-meter *meter-id*

no multicast-meter

Context

config>qos>fc-meter-map>fc

Platforms

7210 SAS-Mxp

Description

This command overrides the default multicast forwarding type meter mapping for the **fc** *fc-name* command. The specified *meter-id* parameter must exist within the policy before the mapping can be made. After the forwarding class mapping is executed, all multicast traffic on a SAP using this policy is forwarded using the meter-id.

The multicast forwarding type includes the unknown **unicast** forwarding type and the **broadcast** forwarding type unless each is explicitly defined to a different meter. When the unknown and broadcast forwarding types are left as default, they track the defined meter for the multicast forwarding type. For more information, see [Use of the keyword "multipoint" for default meter "11"](#) and [Service ingress meter selection rules](#).

The **no** form of this command reverts the multicast forwarding type *meter-id* to the default meter for the forwarding class. If the **broadcast** and **unknown** forwarding types are not explicitly defined to a multipoint meter, they are also reverted to the default multipoint meter. By default, the multicast forwarding type uses meter 1, unless the default multipoint meter 11 is configured.

Parameters

<i>meter-id</i>	Specifies the meter ID for the multicast meter.
Values	1 to 32
Default	1

unknown-meter

Syntax

unknown-meter *meter-id*
no unknown-meter

Context

config>qos>fc-meter-map>fc

Platforms

7210 SAS-Mxp

Description

This command overrides the default unknown unicast forwarding type meter mapping for the **fc** *fc-name* command. The specified *meter-id* parameter must exist within the policy before the mapping can be made. After the forwarding class mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the *meter-id*.

The unknown forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior. For more information, see [Use of the keyword "multipoint" for default meter "11"](#) and [Service ingress meter selection rules](#).

The **no** form of this command reverts the unknown forwarding type to the default of tracking the multicast forwarding type meter mapping.

Parameters

<i>meter-id</i>	Specifies the meter ID for the unknown meter.
Values	1 to 32
Default	1

meter

Syntax

meter *meter-id* [**create**]

no meter *meter-id*

Context

config>qos>fc-meter-map

Platforms

7210 SAS-Mxp

Description

This command configures a meter for an FC meter map policy.

Parameters

meter-id

Specifies the meter ID for the FC meter map policy.

Values 1 to 32

create

Keyword to create a meter.

9.8.2.2 Operational commands



Note:

The 7210 SAS platform QoS capabilities vary across platforms. In the description the term queue/meter is used and based on the platform capabilities both of them or one of them can be used. The description also mentions the capabilities of the node/platform in certain commands, as applicable.

copy

Syntax

copy sap-ingress *src-pol dst-pol* [**overwrite**]

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command copies existing QoS policy entries for a QoS policy ID to another QoS policy ID.

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

Parameters

sap-ingress src-pol dst-pol

Specifies the source policy ID that the **copy** command will attempt to copy from and the destination policy ID to which the command will copy a duplicate of the policy. This parameter indicates that the source policy ID and the destination policy ID are SAP ingress policy IDs.

Values 1 to 65535

overwrite

Keyword 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.

renum

Syntax

renum *old-entry-id new-entry-id*

Context

config>qos>sap-ingress>ip-criteria

config>qos>sap-ingress>ipv6-criteria

config>qos>sap-ingress>mac-criteria

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command renumbers existing QoS policy criteria entries to properly sequence policy entries.

This can be required in some cases since the 7210 SAS exits when the first match is found and executes the actions in accordance with the accompanying **action** command. This requires that entries be sequenced correctly from most to least explicit.

Parameters

old-entry-id

Specifies the entry ID of the existing QoS policy criteria entry.

Values 1 to 64

new-entry-id

Specifies the entry ID for the new QoS policy criteria entry.

Values 1 to 64

9.8.2.3 Show commands

dot1p-classification

Syntax

```
dot1p-classification [policy-id] association
dot1p-classification [policy-id] [detail]
```

Context

show>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays dot1p classification policy information.

Parameters

policy-id
Displays information about the specific policy ID.

Values	1 to 65535
Default	all dot1p classification policies

association
Displays the policy associations of the dot1p classification policy.

detail
Displays detail information for the dot1p classification policy.

Output

The following output is an example of dot1p classification information, and [Table 64: Output fields: dot1p classification policy](#) describes the output fields.

Sample output

```
*A:Dut-A>config>qos# show qos dot1p-classification 10 association
=====
DOT1P Classification Maps
=====
-----
Dot1P Class Id      : 10
Description         : (Not Specified)
-----
Network Policy Associations
-----
```



```

No Network Policy Associations found.
-----

SAP Ingress Associations
-----
SAP Ingress Id          : 100
=====
*A:Dut-A>config>qos# show qos dot1p-classification 10 detail
=====
DOT1P Classification Maps
=====
Dot1P Class Id      : 10
Description         : (Not Specified)

-----
Dot1P Bit Map      Forwarding Class      Profile
-----
2                  l2                   In
3                  ef                   Out
5                  nc                   None
7                  be                   Ukwn

-----
Network Policy Associations
-----
No Network Policy Associations found.
-----

SAP Ingress Associations
-----
SAP Ingress Id          : 100
=====

```

Table 64: Output fields: dot1p classification policy

Label	Description
Dot1p Classification Maps	
Dot1p Class Id	The dot1p classification identifier which identifies the dot1p classification policy
Description	A text string that helps identify the policy context in the configuration file
Default fc	Specifies the default forwarding class for the policy, which is used for all DSCP values that are not explicitly defined in the policy
Default Profile	Specifies the default profile for the policy, which is used for all dot1p values that are not explicitly defined in the policy
Dot1p Bit Map	The dot1p bit value that maps to the forwarding class and profile

Label	Description
Forwarding Class	The forwarding class that maps to the dot1p bit value
Profile	The profile that maps to the dot1p bit value
Port Attachments	
Port-id	The port identifier of a port that is configured to use the dot1p classification policy
SAP Ingress Associations	
SAP Ingress Id	The policy identifier of a SAP ingress QoS policy to which the dot1p classification policy is associated

dscp-classification

Syntax

dscp-classification [*policy-id*] **association**

dscp-classification [*policy-id*] [**detail**]

Context

show>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays DSCP classification policy information.

Parameters

policy-id

Displays information about the specific policy ID.

Values 1 to 65535

Default all DSCP classification policies

associations

Displays the policy associations of the DSCP classification policy.

detail

Displays detail information for the DSCP classification policy.

Output

The following output is an example of DSCP classification policy information, and [Table 65: Output fields: DSCP classification policy](#) describes the output fields.

Sample output

```
*A:Dut-A# show qos dscp-classification 1 association
=====
DSCP Classification Maps
=====
-----
Dscp Class Id      : 1
Description        : Default DSCP Classification policy
-----
Port Attachments
-----
Port-id : 3/1/1
Port-id : 3/1/2
Port-id : 3/1/3
Port-id : 3/1/4
Port-id : 4/1/1
Port-id : 5/1/2
Port-id : 5/1/3
Port-id : 5/1/4
-----
SAP Ingress Associations
-----
SAP Ingress Id      : 1
SAP Ingress Id      : 17
-----
Override policy attachment
-----
Router Interface     : a
=====

*A:Dut-A# show qos dscp-classification 1 detail
=====
DSCP Classification Maps
=====
-----
Dscp Class Id      : 1
Description        : Default DSCP Classification policy
Default fc         : be
Default Profile     : Out
-----
Dscp Bit Map                Forwarding Class                Profile
-----
No Matching Entries
-----
Port Attachments
-----
Port-id : 3/1/1
Port-id : 3/1/2
Port-id : 3/1/3
Port-id : 3/1/4
Port-id : 4/1/1
Port-id : 5/1/2
Port-id : 5/1/3
Port-id : 5/1/4
```

```

-----
SAP Ingress Associations
-----
SAP Ingress Id      : 1
SAP Ingress Id      : 17
-----

Override policy attachment
-----
Router Interface    : a
=====

```

Table 65: Output fields: DSCP classification policy

Label	Description
DSCP Classification Maps	
Dscp Class Id	The DSCP classification identifier which identifies the DSCP classification policy
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, which is used for all DSCP values that are not explicitly defined in the policy
Default Profile	Specifies the default profile for the policy, which is used for all DSCP values that are not explicitly defined in the policy
Dscp Bit Map	The DSCP bit value that maps to the forwarding class and profile
Forwarding Class	The forwarding class that maps to the DSCP bit value
Profile	The profile that maps to the DSCP bit value
Port Attachments	
Port-id	The port identifier of a port that is configured to use the DSCP classification policy
SAP Ingress Associations	
SAP Ingress Id	The policy identifier of a SAP ingress QoS policy to which the DSCP classification policy is associated
Override policy attachment	
Router Interface	Identifies the IES or VPRN interface in an RVPLS that has an override DSCP classification policy configured

sap-ingress

Syntax

sap-ingress [*policy-id*] [**association** | **match-criteria** | **detail**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

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

associations

Displays the policy associations of the sap-ingress policy.

match-criterion

Displays the match-criterion of the sap-ingress policy.

detail

Displays detail information for the sap-ingress policy.

Output

The following outputs are examples of SAP ingress QoS policy information, and [Table 66: Output fields: SAP-ingress QoS policy](#) describes the output fields:

- [Sample output for 7210 SAS-Mxp \(MAC match criteria\)](#)
- [Sample output for 7210 SAS-Mxp \(IP criteria, DSCP value, and DSCP mask configured\)](#)
- [Sample output for 7210 SAS-R6 and 7210 SAS-R12 \(DSCP value and DSCP mask configured\)](#)
- [Sample output for 7210 SAS-R6 and 7210 SAS-R12 \(DSCP name configured\)](#)

Sample output for 7210 SAS-Mxp (MAC match criteria)

```
*A:NS1543C2102# show qos sap-ingress 4 detail
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (4)
```

```

-----
Policy-id           : 4                Scope           : Template
Default FC         : be
Criteria-type      : MAC                IP-Mac Rule Priority : None
Mac Sub-Criteria   : dot1p             IP Sub-Criteria    : None
IPv6 Enabled       : False
DSCP Class Policy Id : 1
Accounting         : packet-based
Classifiers Allowed : 16                Meters Allowed     : 8
Service Meter Enabled: : False
-----

Cam-based Resource Requirement
-----
Classifiers Reqrđ (VPLS) : 3                Meters Reqrđ (VPLS) : 1
Classifiers Reqrđ (L3 Mc) : 3                Meters Reqrđ (L3 Mc) : 1
Classifiers Reqrđ (EPIPE) : 3                Meters Reqrđ (EPIPE) : 1
-----

Table-based Resource Requirement
-----
Classifiers Reqrđ (VPLS) : 2                Meters Reqrđ (VPLS) : 1
Classifiers Reqrđ (L3 Mc) : 2                Meters Reqrđ (L3 Mc) : 1
Classifiers Reqrđ (EPIPE) : 2                Meters Reqrđ (EPIPE) : 1

Name                : (Not Specified)
Description          : (Not Specified)
-----

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
-----

Meter Mode          CIR Admin CIR Rule PIR Admin PIR Rule CBS Admin MBS Admin
Color Mode          CIR Oper          PIR Oper          CBS Oper  MBS Oper
-----
1      TrTcm1        0              closest  max      closest  def kbits  def kbits
      color-blind 0              400000000 32 kbits  1024 kbits
-----

FC          UCastM      MCastM      BCastM      UnknownM
-----

No FC-Map Entries Found.
-----

Match Criteria
-----

Mac Match Criteria
-----
Entry           : 1                FrameType       : Ethernet *
Description     : (Not Specified)
Src MAC        :
Dst MAC        :
Ethernet-type   : Disabled
FC             : be

Entry           : 2                FrameType       : Ethernet *
Description     : (Not Specified)
Src MAC        :

```

```

Dst MAC          :
Ethernet-type    : Disabled
FC              : nc

SAP Associations
-----
Service-Id       : 1 (Epipe)      Customer-Id       : 1
- SAP : 1/1/4

=====
* indicates that the corresponding row element may have been truncated.
*A:NS1543C2102#

```

Sample output for 7210 SAS-Mxp (IP criteria, DSCP value, and DSCP mask configured)

```

*A:NS1543C2102# show qos sap-ingress 3 detail
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (3)
-----
Policy-id        : 3                Scope           : Template
Default FC       : be
Criteria-type    : IP               IP-Mac Rule Priority : None
Mac Sub-Criteria : None             IP Sub-Criteria     : dscp
IPv6 Enabled     : False            IPv6 Sub-Criteria   : dscp
DSCP Class Policy Id : 1
Accounting       : packet-based
Classifiers Allowed : 2              Meters Allowed     : 1
Service Meter Enabled: : False

-----
Cam-based Resource Requirement
-----
Classifiers Reqr (VPLS) : 2          Meters Reqr (VPLS) : 1
Classifiers Reqr (L3 Mc) : 2         Meters Reqr (L3 Mc) : 1
Classifiers Reqr (EPIPE) : 2         Meters Reqr (EPIPE) : 1

-----
Table-based Resource Requirement
-----
Classifiers Reqr (VPLS) : 2          Meters Reqr (VPLS) : 1
Classifiers Reqr (L3 Mc) : 2         Meters Reqr (L3 Mc) : 1
Classifiers Reqr (EPIPE) : 2         Meters Reqr (EPIPE) : 1

Name              : (Not Specified)
Description       : (Not Specified)
-----
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

-----
Meter Mode      CIR Admin CIR Rule PIR Admin PIR Rule CBS Admin MBS Admin
Color Mode     CIR Oper      PIR Oper      CBS Oper  MBS Oper
-----
1      TrTcm1      0          closest max      closest def kbits  def kbits

```

```

color-blind 0          40000000          32 kbits    1024 kbits
-----
FC          UCastM      MCastM      BCastM      UnknownM
-----
No FC-Map Entries Found.
-----
Match Criteria
-----
IP Match Criteria
-----
Entry              : 1
Description         : (Not Specified)
Source IP           : Undefined
Dest. IP            : Undefined
Source Port         : None          Dest. Port         : None
Protocol            : none          DSCP value/mask    : 4/5
Fragment            : Off           Ip Precedence       : None
FC                  : be            Priority            : Default
-----
IPv6 Match Criteria
-----
No Match Criteria Entries found.
SAP Associations
-----
No Associations Found.
=====
*A:NS1543C2102#

```

Table 66: Output fields: SAP-ingress QoS policy

Label	Description
Policy-Id	The ID that uniquely identifies the policy
Scope	Exclusive — Implies that this policy can only be applied to a single SAP Template — Implies that this policy can be applied to multiple SAPs on the router
Default FC	Specifies the default forwarding class for the policy
Criteria-type	IP — Specifies that an IP criteria-based SAP ingress policy is used to select the appropriate ingress meter and corresponding forwarding class for matched traffic MAC — Specifies that a MAC criteria-based SAP is used to select the appropriate ingress meters and corresponding forwarding class for matched traffic
Sub-Criteria-type	Displays the configured sub-criteria-type
Classifiers Allowed	Indicates the number of classifiers allowed for a service
Meters Allowed	Indicates the number of meters allowed for a service

Label	Description
Description	A text string that helps identify the policy's context in the configuration file
Cam-based Resource Requirement	
Classifiers Reqrđ (VPLS)	The number of CAM-based classification entries required for the VPLS service
Meters Reqrđ (VPLS)	The number of CAM-based meters required for the VPLS service
Classifiers Reqrđ (L3 Mc)	The number of CAM-based classification entries required for the Layer 3 multicast service
Meters Reqrđ (L3 Mc)	The number of CAM-based meters required for the Layer 3 multicast service
Classifiers Reqrđ (EPIPE)	The number of CAM-based classification entries required for the Epipe VLL service
Meters Reqrđ (EPIPE)	The number of CAM-based meters required for the Epipe VLL service
Table-based Resource Requirement	
Classifiers Reqrđ (VPLS)	The number of table-based classification entries required for the VPLS service
Meters Reqrđ (VPLS)	The number of table-based meters required for the VPLS service
Classifiers Reqrđ (L3 Mc)	The number of table-based classification entries required for the Layer 3 multicast service
Meters Reqrđ (L3 Mc)	The number of table-based meters required for the Layer 3 multicast service
Classifiers Reqrđ (EPIPE)	The number of table-based classification entries required for the Epipe VLL service
Meters Reqrđ (EPIPE)	The number of table-based meters required for the Epipe VLL service
Meter	Displays the meter ID
Mode	Specifies the configured mode of the meter (trTcm1 or srTcm)
Color Mode	Specifies the configured color mode of the meter (color-blind or color-aware)
CIR Admin	Specifies the administrative Committed Information Rate (CIR) parameters for the meters

Label	Description
CIR Oper	Specifies the operational Committed Information Rate (CIR) parameters for the meters
CIR Rule	<p>min — The operational CIR for the meters will be equal to or greater than the administrative rate specified using the rate command</p> <p>max — The operational CIR for the meter will be equal to or less than the administrative rate specified using the rate command</p> <p>closest — The operational PIR for the meters will be the rate closest to the rate specified using the rate command without exceeding the operational PIR</p>
PIR Admin	Specifies the administrative Peak Information Rate (PIR) parameters for the meters
PIR Oper	Specifies the operational Peak Information Rate (PIR) parameters for the meters
PIR Rule	<p>min — The operational PIR for the meter will be equal to or greater than the administrative rate specified using the rate command</p> <p>max — The operational PIR for the meters will be equal to or less than the administrative rate specified using the rate command</p> <p>closest — The operational PIR for the meters will be the rate closest to the rate specified using the rate command</p>
CBS Admin CBS Oper	<p>def — Specifies the default CBS value for the meters</p> <p>value — Specifies the value to override the default reserved buffers for the meters</p>
MBS Admin MBS Oper	<p>def — Specifies the default MBS value</p> <p>value — Specifies the value to override the default MBS for the meter</p>
FC	Specifies the forwarding class
UCastM	Specifies the default unicast forwarding type meters mapping
MCastM	Specifies the overrides for the default multicast forwarding type meter mapping
BCastM	Specifies the default broadcast forwarding type meters mapping
UnknownM	Specifies the default unknown unicast forwarding type meters mapping

Label	Description
Match Criteria	
Entry	Indicates the entry ID
DSCP	Specifies a DiffServ Code Point (DSCP) name used for an ingress SAP QoS policy match
DSCP value/mask	Specifies a DiffServ Code Point (DSCP) value and mask used for an ingress SAP QoS policy match
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
Ethernet-type	Specifies an Ethernet type II Ethertype value to be used as a Service Ingress QoS policy match
FC	Specifies the entry's forwarding class
Service Association	
Service-Id	The unique service ID number which identifies the service in the service domain
Customer-Id	Specifies the customer ID which identifies the customer to the service
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP ingress policy is applied

Sample output for 7210 SAS-R6 and 7210 SAS-R12 (DSCP value and DSCP mask configured)

```
*A:7210 SAS> show qos sap-ingress 2 detail

=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (2)
-----
Policy-id           : 2           Scope           : Template
Default FC         : be
Criteria-type      : IP           IP-Mac Rule Priority : None
Mac Sub-Criteria   : None        IP Sub-Criteria    : any
IPv6 Enabled       : False       IPv6 Sub-Criteria   : dscp
Accounting         : packet-based
Classifiers Allowed : 2           Meters Allowed     : 1
Classifiers Reqrđ (VPLS) : 2       Meters Reqrđ (VPLS) : 1
Classifiers Reqrđ (L3 Mc) : 2       Meters Reqrđ (L3 Mc) : 1
Classifiers Reqrđ (EPIPE) : 2       Meters Reqrđ (EPIPE) : 1
Description        : (Not Specified)
```

```

-----
Meter Mode      CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS Admin  MBS Admin
Color Mode     CIR Oper              PIR Oper      CBS Oper    MBS Oper
-----
1      TrTcm1    0          closest  max        closest  def        def
      color-blind 0          200000000
-----

FC          UCastM      MCastM      BCastM      UnknownM
-----
No FC-Map Entries Found.

-----
Match Criteria
-----
IP Match Criteria
-----
Entry              : 1
Description        : (Not Specified)
Source IP          : Undefined      Source Port       : None
Dest. IP          : Undefined      Dest. Port       : None
Protocol          : none          DSCP value/mask   : 4/5
Fragment          : Off          Ip Precedence     : None
FC                : be
-----

IPv6 Match Criteria
-----
No Match Criteria Entries found.

SAP Associations
-----
Service-Id          : 1 (Epipe)      Customer-Id       : 1
- SAP : 1/1/24

```

Sample output for 7210 SAS-R6 and 7210 SAS-R12 (DSCP name configured)

```

*A:NS1117C0020>config>service>epipe>sap$ /show qos sap-ingress 2 detail

=====
QoS Sap Ingress
=====

Sap Ingress Policy (2)
-----
Policy-id          : 2          Scope          : Template
Default FC         : be
Criteria-type      : IP        IP-Mac Rule Priority : None
Mac Sub-Criteria   : None      IP Sub-Criteria    : any
IPv6 Enabled       : False     IPv6 Sub-Criteria   : dscp
Accounting         : packet-based
Classifiers Allowed : 2        Meters Allowed     : 1
Classifiers Reqrd (VPLS) : 2    Meters Reqrd (VPLS) : 1
Classifiers Reqrd (L3 Mc) : 2    Meters Reqrd (L3 Mc) : 1
Classifiers Reqrd (EPIPE) : 2    Meters Reqrd (EPIPE) : 1
Description        : (Not Specified)
-----

Meter Mode      CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS Admin  MBS Admin
Color Mode     CIR Oper              PIR Oper      CBS Oper    MBS Oper

```

```
-----
1      TrTcm1      0          closest      max      closest      def      def
      color-blind 0          20000000    32      512
-----

FC          UCastM      MCastM      BCastM      UnknownM
-----
No FC-Map Entries Found.

-----
Match Criteria
-----
IP Match Criteria
-----
Entry              : 1
Description        : (Not Specified)
Source IP          : Undefined          Source Port          : None
Dest. IP           : Undefined          Dest. Port           : None
Protocol           : none              DSCP                 : cp63
Fragment           : Off               Ip Precedence        : None
FC                 : be
-----

IPv6 Match Criteria
-----
No Match Criteria Entries found.

SAP Associations
-----
Service-Id          : 1 (Epipe)          Customer-Id          : 1
- SAP : 1/1/24
```

fc-meter-map

Syntax

```
fc-meter-map [policy-id] association
fc-meter-map [policy-id] [detail]
```

Context

```
show>qos
```

Platforms

```
7210 SAS-Mxp
```

Description

This command displays information about the FC meter map policy, and shows all the SAP ingress policy IDs that use this map.

Parameters

policy-id

Specifies the policy ID, up to 64 characters.

detail

Displays detailed information about the FC meter map policy.

associations

Displays the policy associations of the FC meter map policy.

Output

The following output is an example of FC meter map information, and [Table 67: Output fields: FC meter map policy](#) describes the output fields.

Sample output

```
*A:Dut-A>config>qos# show qos fc-meter-map 1 detail
=====
QoS Fc Meter Map
=====
-----
Fc Meter Map (1)
-----
Policy-id          : 1                      Counter-mode      : forward-drop-count
-----
Meter Mode        CIR Admin CIR Rule PIR Admin PIR Rule CBS Admin MBS Admin
Color Mode        CIR Oper              PIR Oper              CBS Oper  MBS Oper
-----
1    TrTcm1        0              closest    max          closest    def kbits  def kbits
    color-blind 0              400000000
2    TrTcm1        300             closest    4000         closest    def kbits  def kbits
    color-blind 304              304          32 kbits  1000 kbits
4    TrTcm1        0              closest    max          closest    3030 kbytes def kbits
    color-blind 0              400000000          2959 kbytes 1024 kbits
-----
FC          UCastM      MCastM      BCastM      UnknownM
-----
be          1 (def)      def         def         def
l2          1 (def)      4           def         def
ef          1 (def)      def         2           def
-----
SAP Ingress Associations
-----
SAP Ingress Id      : 10
SAP Ingress Id      : 65536
-----
=====
*A:Dut-A>config>qos# show qos sap-ingress 10 detail
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (10)
-----
Policy-id          : 10                      Scope              : Template
Default FC         : be
Criteria-type       : None                    IP-Mac Rule Priority : None
Mac Sub-Criteria    : None                    IP Sub-Criteria     : None
```

```

IPv6 Enabled      : False
DSCP Class Policy Id : 1          DOT1P Class Policy Id: 1
Accounting        : packet-based
Service Meter Enabled: : True

-----
Service-Meter Resource Details
-----
FC Meter Map ID      : 1
Service Meter Counter Mode: Forward_Drop_cou*
Max Number of Meters : 3
Max Number of Counters : 6

Name                 : (Not Specified)
Description           : (Not Specified)

-----
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

-----
Meter Mode          CIR Admin CIR Rule PIR Admin PIR Rule CBS Admin MBS Admin
Color Mode          CIR Oper          PIR Oper          CBS Oper  MBS Oper
-----
1    TrTcm1         0             closest max         closest def kbits  def kbits
    color-blind 0             400000000 32 kbits  1024 kbits
2    TrTcm1         300          closest 4000         closest def kbits  def kbits
    color-blind 304          304         32 kbits  1000 kbits
4    TrTcm1         0             closest max         closest 3030 kbytes def kbits
    color-blind 0             400000000 2959 kbytes 1024 kbits

-----
FC          UCastM      MCastM      BCastM      UnknownM
-----
be          1 (def)      def         def         def
l2          1 (def)      4           def         def
ef          1 (def)      def         2           def

-----
Match Criteria
-----
No Matching Criteria.

SAP Associations
-----
No Associations Found.

=====
* indicates that the corresponding row element may have been truncated.

```

Table 67: Output fields: FC meter map policy

Label	Description
Policy-Id	The ID that uniquely identifies the policy

Label	Description
Counter-mode	in-out-profile-count — Specifies that in-profile and out-profile packets are counted per meter forward-drop-count — Specifies that forwarded and dropped packets are counted per meter
Meter	Displays the meter ID
Mode	Specifies the configured mode of the meter (trTcm1 or srTcm)
Color Mode	Specifies the color mode of the meter (color-blind or color-aware)
CIR Admin	Specifies the administrative Committed Information Rate (CIR) parameters for the meters
CIR Oper	Specifies the operational Committed Information Rate (CIR) parameters for the meters
CIR Rule	min — Specifies that the operational CIR for the meters is equal to or greater than the administrative rate specified using the rate command max — Specifies that the operational CIR for the meter is equal to or less than the administrative rate specified using the rate command closest — Specifies that the operational PIR for the meters are the rate closest to the rate specified using the rate command, without exceeding the operational PIR
PIR Admin	Specifies the administrative Peak Information Rate (PIR) parameters for the meters
PIR Oper	Specifies the operational PIR parameters for the meter
PIR Rule	min — Specifies that the operational PIR for the meter is equal to or greater than the administrative rate specified using the rate command max — Specifies that the operational PIR for the meters is equal to or less than the administrative rate specified using the rate command closest — Specifies that the operational PIR for the meters is the rate closest to the rate specified using the rate command
CBS Admin CBS Oper	def — Specifies the default CBS value for the meters value — Specifies the value to override the default reserved buffers for the meters
MBS Admin MBS Oper	def — Specifies the default MBS value

Label	Description
	value — Specifies the value to override the default MBS for the meter
FC	Specifies the forwarding class
UCastM	Specifies the default unicast forwarding type meter mapping
MCastM	Specifies the overrides for the default multicast forwarding type meter mapping
BCastM	Specifies the default broadcast forwarding type meter mapping
UnknownM	Specifies the default unknown unicast forwarding type meter mapping
Service-Meter Resource Details	
FC Meter Map ID	The unique FC meter map ID number
Service Meter Counter Mode	Specifies the service meter counter mode, either in-out-profile-count or forward-drop-count
Max Number of Meters	Specifies the maximum number of meters
Max Number of Counters	Specifies the maximum number of counters
Name	Specifies the name of the service meter
Description	Provides a description of the service meter

10 Access egress QoS policies on 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE

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

10.1 Overview

An access egress policy defines the queuing for the traffic egressing on the access ports. Access-egress queue policies are used at the Ethernet port and define the bandwidth distribution for the various FC/queue traffic egressing on the Ethernet port.

There is one default access egress policy which is identified as policy ID 1. Each policy has 8 queues available. The Forwarding Class to queue mapping is predefined and cannot be changed. The queue parameters like CIR, PIR, and so on, can be modified. The default policy can be copied but they cannot be deleted or modified.

10.1.1 Basic configurations

A basic access egress QoS policy must conform to the following:

- have a unique access egress QoS policy ID
- have a QoS policy scope of template or exclusive
- queue parameters can be modified, but not deleted

Output example

The following is a sample configuration output for 7210 SAS-T in access-uplink mode.

```
*7210SAS>config>qos>access-egress$ info detail
-----
no description
remark
scope template
queue 1
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 2
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 3
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 4
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
```

```
queue 5
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
queue 6
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
queue 7
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
queue 8
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
fc af create
  no de-mark
  no dotlp
  dotlp-in-profile 7
  dotlp-out-profile 4
exit
-----
```

Output example

The following is a sample configuration output for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE in network mode.

```
*A:01154300003>config>qos>access-egress# info detail
-----
no description
remarking
remark 200
scope template
queue 1
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
queue 2
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
queue 3
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
queue 4
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
queue 5
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
queue 6
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
queue 7
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
```

```
queue 8
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
exit
-----
```

10.1.1.1 Modifying access egress QoS queues

To modify access egress queue parameters specify the following:

- **queue ID value**

On 7210 SAS-T, 8 Queues are identified and are mapped as defined in [Table 31: Forwarding class to queue-ID map](#).

On 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, a queue ID in the policy identifies a FC as per the table 32. [Table 32: Forwarding class-to-queue ID map for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#). In other words, the rate parameters configured for the FC/queue ID applies to the aggregate rate across both unicast and multicast queue for the FC.

- **queue parameters**

Egress queues support configuration of CIR and PIR rates.

Output example

The following is a sample access egress QoS policy configuration output on 7210 SAS-T access-uplink mode.

```
*A:SAS-T>config>qos>access-egress>queue$ info
-----
      adaptation-rule cir max
      rate cir 2000 pir 30000
-----
```

10.1.1.2 Applying access egress QoS policies

Apply access egress policies to the following entities:

- Ethernet ports

A policy can be applied to the ports that are in access mode.

10.1.1.2.1 Ethernet ports

Use the following syntax to apply a access-egress policy to an Ethernet port.

```
config>port#
  ethernet access egress
  qos access-egress-policy-id
```

Output example

The following is a sample port configuration output.

```
*A:card-1>config>port# info
```

```

-----
        shutdown
        ethernet
        access
        egress
        qos 30
        exit
        exit
        exit
-----
*A:card-1>config>port#

```

10.1.1.3 Default access egress QoS policy values

Output example

The following are sample default policy parameters.

```

*A:card-1>config>qos>access-egress# info detail
-----
description "Default Access egress QoS policy."
no remarking
scope template
queue 1
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 2
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 3
    adaptation-rule cir closest pir closest
    rate 0 pir max
exit
queue 4
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 5
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 6
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 7
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 8
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
-----
*A:card-1>config>qos>access-egress#

```

The following table lists the default forwarding class marking values when remarking is enabled on the access egress policy for 7210 SAS devices configured in network mode and access-uplink mode:

Table 68: Default FC marking values for 7210 SAS-T (access-uplink mode)

Default FC value	Network mode	Access-uplink mode
af	dot1p-in-profile 2 dot1p-out-profile 2 dscp-in-profile af11 dscp-out-profile af12	dot1p-in-profile 2 dot1p-out-profile 2
be	dot1p-in-profile 0 dot1p-out-profile 0 dscp-in-profile be dscp-out-profile be	dot1p-in-profile 0 dot1p-out-profile 0
ef	dot1p-in-profile 5 dot1p-out-profile 5 dscp-in-profile ef dscp-out-profile ef	dot1p-in-profile 5 dot1p-out-profile 5
h1	dot1p-in-profile 6 dot1p-out-profile 6 dscp-in-profile nc1 dscp-out-profile nc1	dot1p-in-profile 6 dot1p-out-profile 6
h2	dot1p-in-profile 4 dot1p-out-profile 4 dscp-in-profile af41 dscp-out-profile af41	dot1p-in-profile 4 dot1p-out-profile 4
l1	dot1p-in-profile 3 dot1p-out-profile 3 dscp-in-profile af21 dscp-out-profile af22	dot1p-in-profile 3 dot1p-out-profile 3
l2	dot1p-in-profile 1 dot1p-out-profile 1 dscp-in-profile cs1 dscp-out-profile cs1	dot1p-in-profile 1 dot1p-out-profile 1
nc	dot1p-in-profile 7 dot1p-out-profile 7 dscp-in-profile nc2	dot1p-in-profile 7 dot1p-out-profile 7

Default FC value	Network mode	Access-uplink mode
	dscp-out-profile nc2	

10.1.1.4 Deleting QoS policies

Every access Ethernet port is associated, by default, with the default access egress policy (policy-id 1). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the port configuration. When you remove a non-default access egress policy, the association reverts to the default policy-id 1.

A QoS policy cannot be deleted until it is removed from all access ports where they are applied.

```
*A:card-1>config>qos# no access-egress 30
MINOR: CLI Could not remove Access egress policy "30" because it is in use.
```

10.1.1.5 Removing a policy from the QoS configuration

```
config>qos# no access-egress policy-id
```

Example:

```
config>qos# no access-egress 100
config>qos# no access-egress 1010
```

10.2 Access egress QoS policy command reference

10.2.1 Command hierarchies

- [Configuration commands for 7210 SAS-T \(access-uplink mode\)](#)
- [Configuration commands for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE \(network mode\)](#)
- [Operational commands](#)
- [Show commands](#)

10.2.1.1 Configuration commands for 7210 SAS-T (access-uplink mode)

```
config
- qos
  - access-egress policy-id [create]
  - no access-egress policy-id
    - description description-string
    - no description
    - fc fc-name [create]
    - no fc fc-name
```

```
- [no] de-mark [force de-value]
- dot1p dot1p-priority
- no dot1p
- dot1p-in-profile dot1p-value
- no dot1p-in-profile
- dot1p-out-profile dot1p-value
- no dot1p-out-profile
- dscp-in-profile dscp-name
- no dscp-in-profile
- dscp-out-profile dscp-name
- no dscp-out-profile
- queue queue-id
- adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
- no adaptation-rule
- rate cir cir-rate [pir pir-rate]
- no rate
- remarking {use-dot1p | use-dscp | all}
- no remarking
- scope {exclusive | template}
- no scope
```

10.2.1.2 Configuration commands for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE (network mode)

```
config
- qos
  - access-egress policy-id [create]
  - no access-egress policy-id
  - description description-string
  - no description
  - queue queue-id
    - adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
    - no adaptation-rule
    - rate cir cir-rate [pir pir-rate]
    - no rate
  - remark policy-id
  - no remark
  - remarking
  - no remarking
  - scope {exclusive | template}
  - no scope
```

10.2.1.3 Operational commands

```
- config
  - qos
    - copy sap-ingress src-pol dst-pol overwrite
```

10.2.1.4 Show commands

```
show
- qos
  - access-egress [policy-id] [association| detail]
```


10.2.2 Configuration commands

10.2.2.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>access-egress

Platforms

7210 SAS-T (in access-uplink and network modes), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

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

description-string

Specifies 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.

access-egress

Syntax

access-egress *policy-id* [**create**]

no access-egress *policy-id*

Context

config>qos

Platforms

7210 SAS-T (in access-uplink and network modes), 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-Sx 10/100GE

Description

This command creates or edits an access egress QoS policy. The policy can be applied to multiple access ports. The access egress policy is common to services (SAPs) that are all egressing on a particular port.

Any changes made to an existing policy are applied to all access ports on which the policy is specified.

The remarking parameters and queue parameters are used when port-based queuing is configured

This command is used to create or edit a access egress QoS policy. The egress policy defines the queue parameters (CIR/PIR) for each of the forwarding class traffic as they egress on the access port. Policies in effect are templates that can be applied to multiple access ports as long as the scope of the policy is template. There are 8 queues always available per port for which parameters are configurable.

Parameters

policy-id

Specifies the value that uniquely identifies the access-egress policy.

Values 1 to 65535

create

Keyword to create an access-egress policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

fc

Syntax

fc *fc-name* [**create**]

no fc *fc-name*

Context

config>qos>access-egress

Platforms

7210 SAS-T (access-uplink mode)

Description

This command defines the **fc** node within the access egress QoS policy is used to contain the explicitly defined dot1p marking commands for the *fc-name*.



Note:

When the mapping for the *fc-name* and marking value is not defined, the node for *fc-name* is not displayed in the show configuration or save configuration output.

The **no** form of this command removes the explicit dot1p marking commands for the *fc-name*.

Parameters

fc-name

Specifies the forwarding class for which 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

create

Keyword to create an access-egress policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

de-mark

Syntax

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

Context

config>qos>access-egress>fc

Platforms

7210 SAS-T (access-uplink mode)

Description

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

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

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

Parameters

de-value

Specifies the DEI value to use for this forwarding class.

Values 0 or 1

dot1p

Syntax

[no] **dot1p** *dot1p-value*

Context

config>qos>access-egress>fc

Platforms

7210 SAS-T (access-uplink mode)

Description

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the **dot1p** command has no effect.

DEI marking can be enabled using the **de-mark** command along with this command for the command to take effect. When the **de-mark** command is configured along with this command, the DEI bit is marked in the packet to indicate the profile of the packet. The DEI bit is marked to 0 to indicate in-profile/green packet and 1 to indicate out-of-profile/yellow packet. If the **force de-value** parameter is specified then the DEI bit is set to specified value for all packets.

If the **no** form of this command is executed then software will use the **dot1p-in-profile** and **dot1p-out-profile** if configured, else it will use default values.



Note:

The following rules are applied by software to determine the Dot1p values when the **dot1p**, **dot1p-in-profile**, and **dot1p-out-profile** commands are specified:

1. If **de-mark** is not configured, then dot1p [in | out]-profile values are considered. Even if **dot1p value** is configured, it is configured it is ignored and if 'dot1p [in | out]-profile' value is not configured then default values are considered for that FC.
2. If **de-mark** is configured and if **dot1p value** is configured then it is considered. Else if 'dot1p [in | out]-profile' value is configured it is considered. In this case, **dot1p value**, has the precedence over 'dot1p [in | out]-profile'.

Default

no dot1p

Parameters

dot1p-value

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

dot1p-in-profile

Syntax

dot1p-in-profile *dot1p-priority*

no dot1p-in-profile

Context

config>qos>>access-egress>fc

Platforms

7210 SAS-T (access-uplink mode)

Description

The command will add the capability to mark on an egress the in and out of profile status via a certain dot1p combination, similarly with the DEI options. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with in-profile status (or green color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the dot1p command has no effect.

If DEI marking is enabled using the **de-mark** command and the command 'dot1p *dot1p-value*' is used to configure the dot1p value, then this command has no effect. In other words, enabling DEI marking has precedence over this command and the system ignores this command.

When this command is used the DEI Bit is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DEI bit is set to 0.

The **no** form of this command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.



Note:

The following rules are applied by software to determine the Dot1p values when the **dot1p**, **dot1p-in-profile**, and **dot1p-out-profile** commands are specified:

1. If **de-mark** is not configured, then dot1p [in|out]-profile values are considered. Even if 'dot1p <val>' command is configured it is ignored and if 'dot1p [in|out]-profile' value is not configured then default values are considered for that FC.
2. If **de-mark** is configured and if 'dot1p <value>' command is configured then it is considered. Else if 'dot1p [in|out]-profile' value is configured it is considered. In this case 'dot1p <val>' has the precedence over 'dot1p [in|out]-profile'.
3. If marking is enabled and both dot1p <val> and dot1-[in|out]-profile commands are not specified, then the default values for default.

Default

0

Parameters

dot1p-priority

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

dot1p-out-profile

Syntax

dot1p-out-profile *dot1p-priority*

no dot1p-out-profile

Context

config>qos>access-egress>fc

Platforms

7210 SAS-T (access-uplink mode)

Description

The command will add the capability to mark on an egress the in and out of profile status via a certain dot1p combination, similarly with the DEI options. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with out-of-profile status (or yellow color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the dot1p command has no effect.

If DEI marking is enabled using the **de-mark** command and the *dot1p-value* is configured, then this command has no effect. In other words, enabling DEI marking has precedence over this command and the system ignores this command.

When this command is used the DEI Bit is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DEI bit is set to 0.

The **no** form of this command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.



Note:

The following rules are applied by software to determine the Dot1p values when the **dot1p**, **dot1p-in-profile**, and **dot1p-out-profile** commands are specified:

1. If **de-mark** is not configured, then dot1p [in|out]-profile values are considered. Even if `dot1p <val>' command is configured it is ignored and if `dot1p [in|out]-profile' value is not configured then default values are considered for that FC.
2. If **de-mark** is configured and if `dot1p <value>' command is configured then it is considered. Else if `dot1p [in|out]-profile' value is configured it is considered. In this case `dot1p <val>', has the precedence over `dot1p [in|out]-profile'.
3. If marking is enabled and both dot1p <val> and dot1-[in|out]-profile commands are not specified, then the default values for default.

Default

0

Parameters

dot1p-priority

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

dscp-out-profile

Syntax

dscp-out-profile *dscp-name*

no dscp-out-profile

Context

config>qos>access-egress>fc

Platforms

7210 SAS-T (in access-uplink mode and network).

Description

This command specifies the out-of-profile DSCP name for the forwarding class. When marking is set, the corresponding DSCP value is used to mark all IP packets with out-of-profile status, on the egress of this forwarding class queue.

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

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

Parameters

dscp-name

Specifies the DSCP 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, cp35, 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
---------------	---

dscp-in-profile

Syntax

dscp-in-profile *dscp-name*

no dscp-in-profile

Context

config>qos>access-egress>fc

Platforms

7210 SAS-T (access-uplink mode)

Description

This command specifies the in-profile DSCP name for the forwarding class. When marking is set, the corresponding DSCP value is used to mark all IP packets with out-of-profile status, on the egress of this forwarding class queue.

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

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

Parameters

dscp-name

Specifies the DSCP 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, cp35, 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

queue

Syntax

queue *queue-id*

Context

config>qos>access-egress

Platforms

7210 SAS-T (in access-uplink and network modes), 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-Sx 10/100GE

Description

This command creates the context to modify queue parameters associated with a particular queue. The queue is identifiable by queue ID and FCs are mapped into the queues. See [Table 31: Forwarding class to queue-ID map](#) and [Table 32: Forwarding class-to-queue ID map for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#) for more information.

Parameters

queue-id

Specifies the access egress *queue-id* associated with an FC. See [Table 31: Forwarding class to queue-ID map](#) and [Table 32: Forwarding class-to-queue ID map for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#) for more information.

Values 1 to 8

remarking

Syntax

remarking {**use-dot1p**|**use-dscp**|**all**}

no remarking

Context

config>qos>network>egress

Platforms

7210 SAS-T (access-uplink mode)

Description

This command enables the system to remark egress packets sent out of access ports. The user can specify if dot1p or dscp or both dot1p and dscp to be used for marking the packets sent out of the port.

When 7210 is operated in access-uplink mode, marking support is available as given below:

- On access port egress, the behavior is as follows:
 - If the **use-dot1p** keyword is configured in the access-egress policy, then the dot1p bits are marked in the packet header for all traffic sent out of all SAPs configured on that access port.
 - If the **use-dscp** keyword is configured in the access-egress policy, then the IP DSCP bits are marked in the packet header for IPv4 traffic sent out of all SAPs configured on that access port.



Note:

DSCP marking also marks the IPv4 packets associated with SAPs configured in an Layer 2 VPN service. If this is not required, to avoid this it is recommended to use only dot1p marking on access ports.

- If the **all** keyword is configured in the access-egress policy, then the Dot1p bits are marked in the packet header for all traffic (Layer 2 and IPv4) sent out of all SAPs and the IP DSCP bits are marked in the packet header for all IPv4 traffic sent out of all SAPs configured on that access port.



Note:

- DSCP marking also marks the packets associated with SAPs configured in an Layer 2 VPN service. If this is not required, to avoid this it is recommended to use only dot1p marking on access ports.
- If remarking is enabled in access-egress policy, by default **use-dot1p** is used. If no marking values are specified, then the default FC to Dot1p marking values are used.

The **no** form of this command disables remarking.

Default

no remarking

remark

Syntax

remark *policy-id*

no remark

Context

config>qos>access-egress

Platforms

7210 SAS-T (network mode), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE network mode

Description

This command specifies the remarking policy ID to use for marking packets on access egress (also known as, access port egress).

The remarking policy ID must be associated with the appropriate access egress policy and remarking must be enabled in the access egress policy to enable marking of packets sent out of all SAPs configured on the access port.

Only remarking policy of type dot1p, dscp, or dot1p-dscp is allowed to be used when the remark policy is associated with access-egress. See [Table 78: Summary of remark policy and attachment points for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE \(network mode\)](#) and [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#), for different remark policies supported on the node and its use.

The **no** form of this command removes the explicit association of the remark policy and associates the default remark policy. In other words, if remarking is enabled and no remark policy is executed, then the default remark policy is used to mark packets sent out. If no remark policy is executed and remarking is disabled, packets are not remarked.

Parameters

policy-id

Specifies the remark policy.

Values 1 to 65535

remarking

Syntax

no remarking

remarking

Context

config>qos>access-egress

Platforms

7210 SAS-T (network mode), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE network mode

Description

This command enables the system to remark egress packets sent out of access ports.

When remarking is enabled, the remark policy configured in the QoS policy context is used to determine the FC to QoS bit mapping. For example, when remarking is enabled in the access-egress QoS policy, the remark policy associated with access-egress QoS policy is used to determine the FC to dot1p mapping to use for marking packets sent out of access ports.

See [Remark policies](#) for the remark policy that can be used to configure FC to priority bit markings in different QoS policies associated with different service entities. For more information, see [Table 78: Summary of remark policy and attachment points for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE \(network mode\)](#) and [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#)

The **no** form of this command disables remarking.

Default

no remarking

10.2.2.2 Access egress queue QoS policy commands

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]

no adaptation-rule

Context

config>qos>access-egress>queue

Platforms

7210 SAS-T (in access-uplink and network modes), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command defines the method used by the system to derive the operational CIR and PIR rates when the queue is provisioned in hardware. For the **cir** and **pir** parameters, 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 CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **cir** and **pir** apply.

Default

adaptation-rule pir closest cir closest

Parameters

cir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced when adapting the CIR rate defined using the **queue queue-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used to derive the operational CIR rate for the queue. When the **cir** parameter is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

Values **max** — Specifies that the operational CIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational CIR value is greater than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

pir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced when adapting the PIR rate defined using the **queue queue-id rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR rate for the queue. When the **pir** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

Values **max** — Specifies that the operational PIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational PIR value is greater than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational PIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

rate

Syntax

rate *cir* *cir-rate* [*pir* *pir-rate*]

no rate

Context

config>qos>access-egress>queue

Platforms

7210 SAS-T (in access-uplink and network modes), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the 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.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created on the access ports.

The **no** form of this command reverts all queues created with the queue ID by association with the QoS policy to the default PIR (max) and CIR (0) parameters.

Parameters

cir *cir-rate*

Specifies the administrative CIR rate, in kilobits, for the queue. The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a valid CIR setting must be explicitly defined. 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 1000000, max (For devices with only 1G ports)
 0 to 10000000, max (For devices with both 1G and 10G ports)

Default 0

pir pir-rate

Specifies the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a PIR setting is optional. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

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

Values 0 to 1000000, max (For devices with only 1G ports)
 0 to 10000000, max (For devices with both 1G and 10G ports)

Default max

scope

Syntax

scope {exclusive | template}

no scope

Context

config>qos>access-egress

Platforms

7210 SAS-T (in access-uplink and network modes), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command configures the scope as exclusive or template. The policy's scope cannot be changed if the policy is applied to multiple ports.

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

Default

template

Parameters

exclusive

Specifies that the policy can only be applied to one interface port. If a policy with an exclusive scope is assigned to a second interface an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface. The system default policies cannot be put into the **exclusive** scope. An error will be generated if **scope exclusive** is executed in any policies with a policy ID equal to 1.

template

Specifies that the policy can be applied to multiple ports on the router.

10.2.2.3 Operational commands

copy

Syntax

copy sap-ingress *src-pol dst-pol* [**overwrite**]

Context

config>qos

Platforms

7210 SAS-T (in access-uplink and network modes), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, 7210 SAS-Sx 10/100GE

Description

This command copies existing QoS policy entries for a QoS policy ID to another QoS policy ID.

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

Parameters

sap-ingress *src-pol dst-pol*

Specifies 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. Indicates that the source policy ID and the destination policy ID are SAP ingress policy IDs.

Values 1 to 65535

overwrite

Keyword 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.

10.2.2.4 Show commands

access-egress

Syntax

access-egress [*policy-id*] [**association** | **detail**]

Context

show>qos

Platforms

7210 SAS-T (in access-uplink and network modes), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command displays access egress QoS policy information.

Parameters

policy-id

Displays the policy ID of the access-egress policy.

association

Displays associations related to the specified access-egress policy.

detail

Displays detailed policy information including the policy associations.

Output

The following output is an example of access egress QoS policy information, and [Table 69: Output fields: QoS access egress](#) describes the output fields.

Sample output

```
A:Dut-A>show>qos# access-egress

=====
Access Egress Policies
=====
Policy-Id      Scope      Description
-----
1              Template  Default Access egress QoS policy.
=====

A:Dut-A>show>qos#
A:Dut-A>show>qos# access-egress 1 detail

=====
QoS Access Egress
=====
-----
Policy-id      : 1                      Scope      : Template
Remark        : False                  Remark Pol Id: 2
Description    : Default Access egress QoS policy.
-----

Queue Rates and Rules
-----
-----
QueueId        CIR          CIR Adpt Rule      PIR          PIR Adpt Rule
-----
Queue1         0             closest            max           closest
Queue2         0             closest            max           closest
Queue3         0             closest            max           closest
Queue4         0             closest            max           closest
Queue5         0             closest            max           closest
```


Queue6	0	closest	max	closest
Queue7	0	closest	max	closest
Queue8	0	closest	max	closest

Queue Mode and Weight Details				

QueueId	Mode	Weight		

Queue1	weighted	1		
Queue2	weighted	1		
Queue3	weighted	1		
Queue4	weighted	1		
Queue5	weighted	1		
Queue6	weighted	1		
Queue7	weighted	1		
Queue8	weighted	1		

High Slope				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	70	90	75
Queue2	Down	70	90	75
Queue3	Down	70	90	75
Queue4	Down	70	90	75
Queue5	Down	70	90	75
Queue6	Down	70	90	75
Queue7	Down	70	90	75
Queue8	Down	70	90	75

Low Slope				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	50	75	75
Queue2	Down	50	75	75
Queue3	Down	50	75	75
Queue4	Down	50	75	75
Queue5	Down	50	75	75
Queue6	Down	50	75	75
Queue7	Down	50	75	75
Queue8	Down	50	75	75

Burst Sizes and Time Average Factor				

QueueId	CBS	MBS	Time Average Factor	Queue-Mgmt

Queue1	def	def	7	default
Queue2	def	def	7	default
Queue3	def	def	7	default
Queue4	def	def	7	default
Queue5	def	def	7	default
Queue6	def	def	7	default
Queue7	def	def	7	default
Queue8	def	def	7	default

Associations				

Port-id : 1/1/21				

```
Port-id : 1/1/23
Port-id : 1/1/25
Port-id : 1/1/26
=====
A:Dut-A>show>qos#
A:Dut-A>show>qos#
```

Table 69: Output fields: QoS access egress

Label	Description
Policy-Id	The ID that uniquely identifies the policy
Remark	True — Remarking is enabled for all the dot1q-tagged packets that egress the ports on which the sap-egress QoS policy is applied and remarking is enabled False — Remarking is disabled for the policy
Remark Pol Id	Displays the policy id of the remarking policy
Scope	Exclusive — Implies that this policy can be applied only to a single access egress port Template — Implies that this policy can be applied to multiple access ports on the router
Description	A text string that helps identify the policy's context in the configuration file
Queue Rates and Rules	
QueueId	Displays the Queue identifier associated with the sap-egress QoS policy
CIR	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 Adpt 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 PIR
PIR	Specifies the administrative Peak Information Rate (PIR) parameters for the queue

Label	Description
	The PIR defines the maximum rate that the queue can transmit packets through the access port.
PIR Adpt Rule	<p>min — The operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command</p> <p>max — The operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command</p> <p>closest — The operational PIR for the queue will be the rate closest to the rate specified using the rate command</p>
High Slope/Low slope	
QueueId	Displays the Queue identifier associated with the sap-egress QoS policy
State	Displays the state of the queue. The state of the queue can be either "Up" or "Down"
Start Avg	Specifies the low priority or high priority RED slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero
Max Avg	<p>Specifies the percentage of the shared buffer space for the buffer pool at which point the drop probability becomes "1"</p> <p>This parameter is expressed as a decimal integer.</p>
Max Prob	Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one
Burst Sizes and Time Average Factor	
QueueId	Displays the Queue identifier associated with the sap-egress QoS policy
CBS	Displays the configured CBS value
MBS	Displays the configured MBS value
Time Average Factor	Displays the value of the time average factor in use

11 Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

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

11.1 Overview

On 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, the users have an option to use either port-based egress queuing and shaping or SAP-based egress queuing and shaping for SAPs configured on access ports or hybrid ports. Use the **configure system resource-profile qos port-scheduler-mode** command on the 7210 SAS-Mxp and the **configure system global-res-profile qos port-scheduler-mode** command on the 7210 SAS-R6 and 7210 SAS-R12 to select the mode to be used for SAPs configured on all the ports of the node (in other words, this is a per node setting).

On 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, an access egress policy allows the user to define the marking values for the traffic sent out of the access ports toward the customer. Access egress QoS policies map forwarding class flows to marking values to use. In addition, based on the queuing mode used on access egress, it also defines the per port queue parameters.

11.1.1 Access egress QoS policy for SAP-based queuing mode

The 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 support SAP-based marking for access SAPs and port-based egress marking on access ports. SAP-based marking is only supported for Layer 2 SAPs, that is, SAPs configured in Epipe and VPLS service. If user enables remarking in the SAP egress policy attached to the SAP, the remark policy configured is used to mark the packets sent out of the SAP. If remarking is disabled in the SAP egress policy attached to the SAP, remark policy configured under the access-egress policy associated with the egress access port is used to mark all packets sent out of the Layer 2 SAP configured on the access port. This is known as port-based marking.

Port-based marking is supported primarily for Layer 3 SAPs (that is, SAPs configured in VPRN services and IES services). In other words, SAP-based marking is not supported for Layer 3 SAPs.

On 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, no explicit CLI command is provided to choose between port-based marking and SAP-based marking for Layer 2 SAPs. The user can choose SAP-based marking by enabling remarking in the SAP egress policy attached to the Layer 2 SAP or choose port-based marking by disabling remarking in the SAP egress policy attached to the SAP and enabling remarking in the access-egress policy associated with the access port on which the Layer 2 SAP is configured.

A remarking policy can be defined for each access egress policy and remarking is disabled by default. Only remarking policy of type dot1p, dot1p-lsp-exp-shared, dscp or dot1p-dscp can be used with access-egress policy. The following is the marking behavior with different remark policy types:

- If remark policy type is dot1p or dot1p-lsp-exp-shared, then all traffic sent out of Layer 2 SAPs and Layer 3 SAPs configured on that port will have its Dot1p bits marked.
- If remark policy type is dscp, then all traffic sent out of Layer 2 SAPs and Layer 3 SAPs configured on that port will have its IP DSCP bits marked (assuming Layer 2 SAPs are carrying IP traffic).

- If remark policy type is of type dot1p-dscp, then all traffic sent out of Layer 2 SAPs and Layer 3 SAPs configured on that port will have its IP DSCP bits (assuming Layer 2 SAPs are carrying IP traffic) and dot1p bits marked.



Note:

- On the 7210 SAS-Mxp, for Layer 2 SAPs, if remarking is enabled in the SAP egress policy and port-based marking is disabled, the dot1p values configured in the SAP egress policy are used. For Layer 3 SAPs no marking is done.
- On the 7210 SAS-Mxp, Layer 2 SAPs, if remarking is enabled in the SAP egress policy and port-based marking is enabled, the dot1p values configured in the SAP egress policy are used. For Layer 3 SAPs, the dot1p and DSCP values configured in the access-egress policy are used. In addition, the DSCP values configured in the access-egress policy are used to mark the IP traffic sent out of Layer 2 SAPs.
- On the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, if remarking is disabled for the SAP egress policy and port-based marking is enabled, IP DSCP values are marked, including for the traffic egressing from the Layer 2 SAPs configured on the port. To avoid this, it is recommended to use only FC-to-dot1p values when both Layer 2 and Layer 3 SAPs are configured on the same access port.
- On the 7210 SAS-R6 and 7210 SAS-R12, if remarking is enabled for the SAP egress policy and port based marking is enabled, the values configured in the SAP egress policy are used. For L3 SAPs the values configured in the access-egress policy are used.

11.1.2 Access egress QoS policy for port-based queuing mode

In addition to marking values, the access egress QoS policy provides an option to define port-based queues and scheduling when per port queues are used for SAPs configured on access ports.

Users have an option to use either port-based egress queuing and shaping or SAP-based egress queuing and shaping for SAPs configured on access ports or hybrid ports. Use the **configure system resource-profile qos port-scheduler-mode** command on 7210 SAS-Mxp or the **configure system global-resource-profile qos port-scheduler-mode** on the 7210 SAS-R6 and 7210 SAS-R12 to select the mode to be used for SAPs configured on all the ports of the node. When **port-scheduler-mode** is enabled, software uses 8 egress queues per access port and all the SAPs configured on the port will share the 8 egress queues for traffic sent out of that port. In this mode, SAPs configured on hybrid port shares the egress queues with network port traffic. Enabling **port-scheduler-mode** affects the behavior for all the SAPs configured on either access or hybrid port. That is, port-based egress queues is mutually exclusive to use of SAP-based egress queues. When **port-scheduler-mode** is enabled, per port egress queues are defined using the access egress policies.

Additionally, the marking values used to mark traffic from different forwarding classes is defined by the remark policy in the access egress policy. In other words, per SAP marking cannot be used when Port-based queuing mode is used. A remarking policy can be defined for each access egress policy and remarking is disabled by default. Only remarking policy of type **dot1p**, **dot1p-lsp-exp-shared**, **dscp** or **dot1p-dscp** can be used with access-egress policy. The following is the marking behavior with different remark policy types:

- If the remark policy type is **dot1p** or **dot1p-lsp-exp-shared**, all traffic sent out of Layer 2 SAPs and Layer 3 SAPs configured on that port will have its Dot1p bits marked.
- If the remark policy type is **dscp**, all traffic sent out of Layer 2 SAPs and Layer 3 SAPs configured on that port will have its IP DSCP bits marked (assuming Layer 2 SAPs are carrying IP traffic).

- If remark policy is of type **dot1p-dscp**, all traffic sent out of Layer 2 SAPs and Layer 3 SAPs configured on that port will have its IP DSCP bits (assuming Layer 2 SAPs are carrying IP traffic) and Dot1p bits marked.



Note:

- When **port-scheduler-mode** is disabled, per-SAP egress queues are available for use. Per-SAP egress queues are configured in the service egress policies.
- On the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, when port-based queuing is enabled, RVPLS SAPs use the port-based egress queues for both unicast and BUM traffic; all the SAPs, including RVPLS SAPs, share the eight egress queues created per port.

Access egress QoS policies define egress queues and map forwarding class flows to queues, if **port-scheduler-mode** is enabled. In **port-scheduler-mode**, the system allocates 8 queues to access port egress by default. To define a basic access egress QoS policy, the following are required:

- unique access egress QoS policy ID
- QoS policy scope of template or exclusive

The parameters that can be configured for a queue are discussed in [Queue parameters](#).

Optional service egress QoS policy elements include:

- specify remark policy that defines IEEE 802.1p priority value remarking based on forwarding class

When port-based queuing is used, the FC to queue map is fixed and the queues priority is determined by the queue number, with higher queue number having the higher priority. The user can configure a queue to be a strict queue to change the scheduling behavior for that queue.

Output example

The following is a sample configuration output showing a default access-egress policy for the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

```
*A:sasr_dutb>config>qos>access-egress# info detail
-----
description "Default Access egress QoS policy."
no remarking
remark 2
scope template
queue 1
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    queue-mgmt "default"
    queue-mode weighted
    weight 1
exit
A:sasr_dutb>config>qos>access-egress#
```

11.1.3 Access egress QoS policy queue override

The following QoS policy queue parameters can be overridden using the **configure>port>ethernet>access>egress>queue-override** command. For command description details, see the "Port Ethernet QoS commands" section in the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide*:

- **adaptation rule**

Specifies the criteria used to compute the operational PIR and CIR values for this queue (**min**, **max**, or **closest**).

- **percent rate**

Enables the configuration of queue PIR and CIR as a percentage of the egress port line rate rather than the values set by the **rate** parameter.

- **queue management**

Associates a queue management policy with the queue.

- **queue mode**

Sets strict or weighted mode for the queue.

- **rate**

Defines the administrative PIR and CIR for the queue.

- **weight**

Determines the proportion of available bandwidth that is given to this queue.

See "Port Ethernet QoS commands" in the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for more information about the **queue-override** command.

11.1.4 Basic configurations

A basic access egress QoS policy must conform to the following:

- have a unique access egress QoS policy ID
- have a QoS policy scope of template or exclusive

11.1.4.1 Modifying access egress QoS queues on the 7210 SAS-R6 and 7210 SAS-R12

Example

The following is a sample **port-scheduler-mode** policy configuration output for the 7210 SAS-R6 and 7210 SAS-R12.

```
*A:sasr_dutb>config>system>glob-res# info detail
-----
      qos
          no port-scheduler-mode
          exit
-----
*A:sasr_dutb>config>system>glob-res#
```

11.1.4.2 Applying access egress QoS policies on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Apply access egress policies to the following entities:

- Ethernet ports

A policy can be applied to the ports that are in access mode.

11.1.4.2.1 Ethernet ports

Use the following syntax to apply an access-egress policy to an Ethernet port.

```
config>port#  
  ethernet access egress  
    qos access-egress-policy-id
```

Output example

The following is a sample port configuration output for 7210 SAS-R6 and 7210 SAS-R12.

```
*A:Dut-B>config>port# ethernet mode access  
*A:Dut-B>config>port# info detail  
-----  
shutdown  
description "10/100/Gig Ethernet SFP"  
ethernet  
  mode access  
  no enable-table-classification  
  access  
    no accounting-policy  
    no    collect-stats  
    egress  
      qos 1  
  exit  
exit  
encap-type null  
exit
```

11.1.4.3 Editing QoS policies

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors perform the following:

1. Copy the policy to a work area.
2. Edit the policy.
3. Overwrite the original policy.

11.1.4.4 Deleting QoS policies

Every access Ethernet port is associated, by default, with the default access egress policy (policy-id 1). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the port configuration. When you remove a non-default access egress policy, the association reverts to the default policy-id 1.

A QoS policy cannot be deleted until it is removed from all access ports where they are applied.

Example

```
*A:7210-SAS-1>config>qos# no access-egress 30  
MINOR: CLI Could not remove Access egress policy "30" because it is in use.
```


11.1.4.5 Removing a policy from the QoS configuration

```
config>qos# no access-egress policy-id
```

Example:

```
config>qos# no access-egress 100  
config>qos# no access-egress 1010
```

11.2 Access egress QoS policy command reference

11.2.1 Command hierarchies

- [Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#)
- [Show commands](#)

11.2.1.1 Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

```
config  
- qos  
  - access-egress policy-id [create]  
  - no access-egress policy-id  
    - [no] description description-string  
    - queue queue-id  
      - [no] adaptation-rule [cir adaptation-rule] [pir adaptation-rule]  
      - percent-rate cir cir-percent [pir pir-percent ]  
      - no percent-rate  
      - no queue-mgmt  
      - queue-mgmt name  
      - no queue-mgmt  
      - queue-mode queue-mode  
      - no queue-mode  
      - [no] rate cir cir-rate [pir pir-rate]  
      - no rate  
      - weight weight  
      - no weight  
    - scope {exclusive | template}  
    - remark policy-id  
    - no remark  
    - remarking  
    - no remarking
```

11.2.1.2 Show commands

```
show  
- qos  
  - access-egress [policy-id] [association| detail]
```

11.2.2 Command descriptions

11.2.2.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>access-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

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

description-string

Specifies 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.

access-egress

Syntax

access-egress *policy-id* [**create**]

no access-egress *policy-id*

Context

config>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

Commands in this context create or edit an access egress QoS policy. The policy can be applied to multiple access ports. The access egress policy is common to services (SAPs) that are all egressing on a particular port.

On 7210 SAS-Mxp, the access egress policy provides an option to configure remarking parameters and access port egress queue parameters on access port. The remarking parameters are used when SAP-based queuing is configured and port-based marking is in use. The remarking parameters and queue parameters are used when port-based queuing is configured. See [Access egress QoS policies on 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE](#) for more information.

Any changes made to an existing policy are applied to all access ports on which the policy is specified.

A policy cannot be deleted until it is removed from all access ports where it is applied. When an access-egress policy is removed from an access port, the access port reverts to the default access-egress policy ID 1.

The **no** form of the policy associates the default policy with the access port.

Parameters

policy-id

Specifies the value that uniquely identifies the access-egress policy.

Values 1 to 65535

create

Keyword to create an access-egress policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

queue

Syntax

queue *queue-id*

Context

config>qos>access-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

Commands in this context modify queue parameters associated with a particular queue. The queue is identifiable by the *queue-id* and FCs are mapped into the queues. See [Table 31: Forwarding class to queue-ID map](#) for more information.

Parameters

queue-id

Specifies the access egress queue ID associated with an FC.

Values 1 to 8

remark

Syntax

remark *policy-id*

no remark

Context

config>qos>access-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command specifies the remarking policy ID to use for marking packets on access egress (also known as, access port egress).

7210 SAS-Mxp supports two different access egress queuing modes. Marking functionality available in these two modes are different as described in the overview below. See [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information.

SAP-based queuing mode

The 7210 SAS-Mxp supports SAP-based marking for access SAPs and port-based egress marking on access ports. SAP-based marking is only supported for L2 SAPs, that is, SAPs configured in Epipe and VPLS service. If user enables remarking in the SAP egress policy attached to the SAP, then the remark policy configured is used to mark the packets sent out of the SAP. If remarking is disabled in the SAP egress policy attached to the SAP, then remark policy configured under the access-egress policy associated with the egress access port is used to mark all packets sent out of the L2 SAP configured on the access port. This is known as port-based marking. For more details refer to the details in the chapter above on Access Egress policies. section [Access egress QoS policy for SAP-based queuing mode on 7210 SAS-Mxp](#).

Port-based Queuing Mode

When port-based queues are used, only per port egress marking is supported. In other words, per SAP marking cannot be used when Port-based queuing mode is used. For more details refer to the details in the chapter above on Access Egress policies, section [Access egress QoS policy for port-based queuing mode on 7210 SAS-Mxp](#)

The remarking policy ID must be associated with the appropriate access egress policy and remarking must be enabled in the access egress policy to enable marking of packets sent out of all SAPs configured on the access port. Only a remarking policy of the type dot1p, dscp, or dot1p-dscp is allowed to be used when the remark policy is associated with access egress.

See [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information.

The **no** form of this command removes the explicit association of the remark policy and associates the default remark policy. In other words, if remarking is enabled and no remark policy is executed, then the

default remark policy is used to mark packets sent out. If no remark policy is executed and remarking is disabled, then packets are not remarked.

Parameters

policy-id

Specifies the remark policy.

Values 1 to 65535

marking

Syntax

no remarking

marking

Context

config>qos>access-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command enables the system to remark egress packets sent out of access ports.

When remarking is enabled, the remark policy configured in the QoS policy context is used to determine the FC to QoS bit mapping. For example, when remarking is enabled in the access egress QoS policy, the remark policy associated with access-egress QoS policy is used to determine the FC to dot1p mapping to use for marking packets sent out of access ports.

See [Remark policies](#) and [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for the remark policy that can be used to configure FC to priority bit markings in different QoS policies associated with different service entities.



Note:

See [Access egress QoS policies on 7210 SAS-Mxp](#) for more information.

The **no** form of this command disables remarking.

Default

no remarking

11.2.2.2 Access egress queue QoS policy commands

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]
no adaptation-rule

Context

config>qos>access-egress>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command defines the method used by the system to derive the operational CIR and PIR rates when the queue is provisioned in hardware. For the **cir** and **pir** parameters, 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 CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **cir** and **pir** apply.

Default

adaptation-rule pir closest cir closest

Parameters

cir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced when adapting the CIR rate defined using the **queue queue-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used to derive the operational CIR rate for the queue. When the **cir** parameter is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

Values

max — Specifies that the operational CIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational CIR value is greater than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

pir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced when adapting the PIR rate defined using the **queue queue-id rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR rate for the queue. When the **pir** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

Values **max** — Specifies that the operational PIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational PIR value is greater than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational PIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

percent-rate

Syntax

percent-rate [**cir** *cir-percent*] [**pir** *pir-percent*]
no percent-rate

Context

config>qos>access-egress>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command enables support for a queue's PIR and CIR to be configured as a percentage of the egress port's line rate (that is, the port limit). When the rates are expressed as a port limit, the actual rates used per instance of the queue will vary based on the port speed or the configured port egress rate. For example, when the same QoS policy is used on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port because of the difference in port speeds. This enables the same QoS policy to be used on SAPs on different ports instead of using different policies 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 are 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 anytime.

The **no** form of this command reverts the queue to its default shaping rate and CIR values. When **no percent-rate** is defined within an egress queue-override, the queue reverts to the PIR and CIR defined within the access egress QoS policy associated with the queue.

Parameters

cir *cir-percent*

Specifies the queue's CIR as a percentage dependent on the use of the port-limit.

Values 0.00 to 100.00 percent

Default 0.00

pir *pir-percent*

Specifies the queue's PIR as a percentage dependent on the use of the port-limit.

Values 0.01 to 100.00 percent

Default 100.00

queue-mgmt

Syntax

queue-mgmt *name*

no queue-mgmt

Context

config>qos>access-egress>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command associates the specified queue management policy with this queue.

The queue management policy is used to specify the queue buffer parameters and queue slope policy parameters.

The **no** form of this command associates the default SAP egress queue management policy with this queue.

Parameters

name

Specifies the name of the queue management policy, up to 32 characters.

queue-mode

Syntax

queue-mode *queue-mode*

no queue-mode

Context

config>qos>access-egress>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command determines whether the queue operates in strict or weighted mode.

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

Default

weighted

Parameters

queue-mode

Specifies the mode of operation for the queue.

Values **strict** — If a queue is configured in strict mode, the scheduler schedules the queue in order of their priority in two (2) passes, the CIR loop and the PIR loop.

weighted — If a queue is configured in weighted mode, then the scheduler examines these queues in two (2) passes, the CIR loop and the PIR loop. In the CIR loop, it distributes the available bandwidth to all the strict and then weighted queues in round-robin up to the configured CIR rate. It examines the weighted queues in the PIR loop, after examining all the strict queues and distributes the available bandwidth, if any, in the proportion of the configured weights.

rate

Syntax

rate cir *cir-rate* [pir *pir-rate*]

no rate

Context

config>qos>access-egress>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the 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 over-subscription 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.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created on the access ports.

The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the *pir-rate* value.



Note:

Queues with a **cir-level** *cir-level* parameter value of 8 are treated differently by the software than queues configured with different *cir-level* values. The PIR rate values configured for the *cir-level* 8 queues are ignored. Only the CIR rate value is used and PIR is set to the CIR value. In addition, when executing the **no** form of the **rate** command for a queue configured at *cir-level* 8, the default CIR (and PIR) value is set to 1. See the *cir-level* parameter description under the **port-parent** command for more information about the scheduler behavior for queues configured with a *cir-level* value of 8.

The **no** form of this command reverts all queues created with the queue ID by association with the QoS policy to the default PIR (max) and CIR (0) parameters.

Parameters

cir-rate

Specifies the administrative CIR rate, in kilobits, for the queue. The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a valid CIR setting must be explicitly defined. 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 1000000, max (For devices with only 1G ports)
 0 to 10000000, max (For devices with both 1G and 10G ports)

Default 0

pir-rate

Specifies the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a PIR setting is optional. 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 **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 0 to 1000000, max (For devices with only 1G ports)
 0 to 10000000, max (For devices with both 1G and 10G ports)

Default max

weight

Syntax

[no] **weight** *weight*

Context

config>qos>access-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command specifies the weight of the queue when the queue mode is set to **weighted**.

For queues configured in weighted mode, the CIR of the queues are met if bandwidth is available and the configured weights are considered in the PIR loop. The configured weight determines the proportion of available bandwidth that is given to this queue in comparison to other queues configured in weighted mode.

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

Default

1

Parameters

weight

Specifies the proportion of available bandwidth to be allocated to this queue relative to other queues.

Values 1 to 15

scope

Syntax

scope {**exclusive** | **template**}

no scope

Context

config>qos>access-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures the scope as exclusive or template. The policy's scope cannot be changed if the policy is applied to multiple ports.

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

Default

template

Parameters

exclusive

Specifies that the policy can only be applied to one interface port. If a policy with an **exclusive** scope is assigned to a second interface, an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface. The system default policies cannot be put into the **exclusive** scope. An error will be generated if **scope exclusive** is executed in any policies with a policy ID equal to 1.

template

Specifies that the policy can be applied to multiple ports on the router.

11.2.2.3 Show commands

access-egress

Syntax

access-egress [*policy-id*] [**association** | **detail**]

Context

show>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays access egress QoS policy information.

Parameters

- policy-id**
Displays the policy ID of the access-egress policy.
- association**
Displays associations related to the specified access-egress policy.
- detail**
Displays detailed policy information including the policy associations.

Output

The following output is an example of access egress QoS policy information, and [Table 70: Output fields: QoS access egress](#) describes the output fields.

Sample output

```
A:Dut-A>show>qos# access-egress

=====
Access Egress Policies
=====
Policy-Id      Scope      Description
-----
1              Template  Default Access egress QoS policy.
=====

A:Dut-A>show>qos#
A:Dut-A>show>qos# access-egress 1 detail

=====
QoS Access Egress
=====
-----
Policy-id      : 1                      Scope      : Template
Remark        : False                  Remark Pol Id: 2
Description    : Default Access egress QoS policy.
-----

Queue Rates and Rules
-----
-----
QueueId        CIR          CIR Adpt Rule      PIR          PIR Adpt Rule
-----
Queue1         0              closest            max           closest
Queue2         0              closest            max           closest
Queue3         0              closest            max           closest
Queue4         0              closest            max           closest
Queue5         0              closest            max           closest
Queue6         0              closest            max           closest
Queue7         0              closest            max           closest
Queue8         0              closest            max           closest
-----

Queue Mode and Weight Details
-----
-----
```

```

-----
QueueId      Mode      Weight
-----
Queue1       weighted    1
Queue2       weighted    1
Queue3       weighted    1
Queue4       weighted    1
Queue5       weighted    1
Queue6       weighted    1
Queue7       weighted    1
Queue8       weighted    1
-----

High Slope
-----

QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down       70             90           75
Queue2       Down       70             90           75
Queue3       Down       70             90           75
Queue4       Down       70             90           75
Queue5       Down       70             90           75
Queue6       Down       70             90           75
Queue7       Down       70             90           75
Queue8       Down       70             90           75
-----

Low Slope
-----

QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down       50             75           75
Queue2       Down       50             75           75
Queue3       Down       50             75           75
Queue4       Down       50             75           75
Queue5       Down       50             75           75
Queue6       Down       50             75           75
Queue7       Down       50             75           75
Queue8       Down       50             75           75
-----

Burst Sizes and Time Average Factor
-----

QueueId      CBS      MBS      Time Average Factor  Queue-Mgmt
-----
Queue1       def      def      7                    default
Queue2       def      def      7                    default
Queue3       def      def      7                    default
Queue4       def      def      7                    default
Queue5       def      def      7                    default
Queue6       def      def      7                    default
Queue7       def      def      7                    default
Queue8       def      def      7                    default
-----

Associations
-----
Port-id : 1/1/21
Port-id : 1/1/23
Port-id : 1/1/25
Port-id : 1/1/26
=====
A:Dut-A>show>qos#

```

Table 70: Output fields: QoS access egress

Label	Description
Policy-Id	The ID that uniquely identifies the policy
Remark	True — Remarking is enabled for all the Dot1q-tagged packets that egress the ports on which the sap-egress QoS policy is applied and remarking is enabled False — Remarking is disabled for the policy
Remark Pol Id	Displays the policy ID of the remarking policy
Scope	Exclusive — Implies that this policy can be applied only to a single access egress port Template — Implies that this policy can be applied to multiple access ports on the router
Description	A text string that helps identify the policy's context in the configuration file
Queue Rates and Rules	
QueueId	Displays the Queue identifier associated with the sap-egress QoS policy
CIR	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 Adpt 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 PIR
PIR	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 access port.
PIR Adpt Rule	min — The operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command

Label	Description
	<p>max — The operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command</p> <p>closest — The operational PIR for the queue will be the rate closest to the rate specified using the rate command</p>
High Slope/Low slope	
QueueId	Displays the Queue identifier associated with the sap-egress QoS policy
State	<p>Displays the state of the queue</p> <p>The state of the queue can be either "Up" or "Down"</p>
Start Avg	Specifies the low priority or high priority RED slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero
Max Avg	<p>Specifies the percentage of the shared buffer space for the buffer pool at which point the drop probability becomes "1"</p> <p>This parameter is expressed as a decimal integer.</p>
Max Prob	Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one
Burst Sizes and Time Average Factor	
QueueId	Displays the Queue identifier associated with the sap-egress QoS policy
CBS	Displays the configured CBS value
MBS	Displays the configured MBS value
Time Average Factor	Displays the value of the time average factor in use
Queue-Mgmt	Displays the Queue management policy in use

12 Service egress policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

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



Note:

Service egress policies are supported only on the 7210 SAS-Mxp, 7210 SAS-R6 IMM-b, and 7210 SAS-R12 IMM-b. It is not supported on the 7210 SAS-R6 IMM-c, 7210 SAS-R12 IMM-c, 7210 SAS-T, and 7210 SAS-Sx/S 1/10GE devices.

12.1 Overview



Note:

On the 7210 SAS-R6 and 7210 SAS-R12, only unicast traffic sent out of RVPLS SAPs uses per-SAP egress queues. BUM traffic sent out of RVPLS SAPs uses per-port egress queues.

The Service Egress policy defines the Service Level Agreement (SLA) for service packets as they egress on the access SAP. Service Egress QoS policies allow the definition of queue parameters along with a remark policy.

Service Egress Qos policies are available for use only when per SAP egress queues (SAP-based queuing mode) are used. It is not available when port-based queuing mode (that is, with per access port egress queues) is in use.



Note:

The 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, on port egress, use either port-based egress queuing and shaping or SAP-based egress queuing and shaping for all SAPs configured on the port. The **config system global-resource-profile qos port-scheduler-mode** command on the 7210 SAS-R6 and 7210 SAS-R12 and the **config system resource-profile qos port-scheduler-mode** command on the 7210 SAS-Mxp allows the user to select the mode for all SAPs configured on either access ports or hybrid ports. When **port-scheduler-mode** is enabled, the software allocates eight egress queues per port (access port or hybrid port). See [Schedulers on 7210 SAS-R6 and 7210 SAS-R12](#) for more information. When **port-scheduler-mode** is disabled, the software allocates eight egress queues per SAP, which are configured using service egress policies.

12.1.1 Basic configurations

To use per SAP service egress policies, the following must be configured:

- Disable port-scheduler mode using the **config system res-prof qos no port-scheduler-mode** command on the 7210 SAS-Mxp, and **config system global-res-prof no port-scheduler-mode** on the 7210 SAS-R6 and 7210 SAS-R12.

- Ensure resources are available for use by SAP egress queues in the ingress-internal-tcam resources pool using the **config system resource-profile ingress-internal-tcam qos-sap-egress-resource** command.

A basic service egress QoS policy must conform to the following:

- have a unique service egress QoS policy ID
- have a QoS policy scope of template or exclusive
- have at least one forwarding class queue

12.1.2 Create a SAP egress policy

To create a new SAP Egress policy, define the following:

- a SAP egress policy name
- a brief description of the policy features
- the queue parameters for all the queues

Use the following syntax to configure a SAP egress policy.

```
A:Dut-A>config>qos# sap-egress
- no sap-egress <policy-id>
- sap-egress <policy-id> [create]

<policy-id>          : [1..65535]|<name:64 char max>
<create>             : keyword - mandatory while creating an entry.

[no] description      - Description for this sap-egress policy
[no] queue            + Configure a queue
[no] remark           - Specify Remarking policy for this policy
[no] remarking        - Enable/disable remarking
[no] scope            - Specify scope of the policy
```

Output example

The following is a sample SAP Egress policy configuration output for the 7210 SAS-Mxp.

```
*A:sim_dutc>config>qos>sap-egress# info detail
-----
description "Default SAP egress QoS policy."
scope template
no remarking
remark 2
queue 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
    queue-mode weighted
    weight 1
exit
queue 2
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
    queue-mode weighted
```

```
        weight 1
    exit
    queue 3
        adaptation-rule pir closest cir closest
        rate cir 0 pir max
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
    queue 4
        adaptation-rule pir closest cir closest
        rate cir 0 pir max
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
    queue 5
        adaptation-rule pir closest cir closest
        rate cir 0 pir max
        queue-mgmt "default"
        queue-mode weighted
        weight 1
    exit
-----
A:Dut-B>config>qos>sap-egress#
```

Output example

The following is a sample SAP Egress configuration output for the 7210 SAS-R6 and 7210 SAS-R12.

```
*A:SAS>config>qos>sap-egress# info detail
-----
description "Default SAP egress QoS policy."
scope template
no remarking
remark 2
queue 1
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 2
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 3
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 4
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 5
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
```

```
        rate cir 0 pir max
        queue-mgmt "default"
    exit
    queue 6
        port-parent cir-level 1 pir-weight 1
        adaptation-rule pir closest cir closest
        rate cir 0 pir max
        queue-mgmt "default"
    exit
    queue 7
        port-parent cir-level 1 pir-weight 1
        adaptation-rule pir closest cir closest
        rate cir 0 pir max
        queue-mgmt "default"
    exit
    queue 8
        port-parent cir-level 1 pir-weight 1
        adaptation-rule pir closest cir closest
        rate cir 0 pir max
        queue-mgmt "default"
    exit
-----
*A: SAS>config>qos>sap-egress#
```

12.1.3 Editing QoS policies

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors, perform the following:

1. Copy the policy to a work area.
2. Edit the policy.
3. Overwrite the original policy.

12.2 Service egress policy command reference

12.2.1 Command hierarchies

- [Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#)
- [Operational commands](#)
- [Show commands](#)

12.2.1.1 Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

```
config
- qos
- sap-egress policy-id [create]
- no sap-egress policy-id
- [no] description description-string
- queue queue-id
```

```
- [no] adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
- queue-mgmt name
- no queue-mgmt
- queue-mode queue-mode
- no queue-mode
- [no] rate cir cir-rate pir pir-rate
- weight weight
- no weight
- scope {exclusive | template}
- no scope
- [no] remark policy-id
- [no] remarking
```

12.2.1.2 Operational commands

```
config
- qos
- copy sap-egress src-pol dst-pol [overwrite]
```

12.2.1.3 Show commands

```
show
- qos
- sap-egress [policy-id] [detail | association]
```

12.2.2 Command descriptions

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

12.2.2.1 Configuration commands

- [Generic commands](#)
- [SAP egress queue QoS policy commands](#)

12.2.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>sap-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

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

description-string

Specifies 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.

12.2.2.1.2 SAP egress queue QoS policy commands

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]

no adaptation-rule

Context

config>qos>sap-egress>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command defines the method used by the system to derive the operational CIR and PIR rates when the queue is provisioned in hardware. For the **cir** and **pir** parameters, 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 CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **cir** and **pir** apply.

Default

adaptation-rule pir closest cir closest

Parameters

cir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced when adapting the CIR rate defined using the **queue queue-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used to derive the operational CIR rate for the queue. When the **cir** parameter is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) lists the supported hardware step-size rates.

Default closest

Values **max** — Specifies that the operational CIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational CIR value is greater than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

pir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced when adapting the PIR rate defined using the **queue queue-id rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR rate for the queue. When the **pir** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) lists the supported hardware step-size rates.

Default closest

Values **max** — Specifies that the operational PIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational PIR value is greater than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational PIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

queue

Syntax

queue *queue-id* **create**

Context

config>qos>sap-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures queue parameters.

Parameters

queue-id

Specifies the ID of the queue.

Values 1 to 8

create

Keyword to create a network queue policy.

sap-egress

Syntax

sap-egress *policy-id* [**create**]

Context

config>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures a SAP Egress policy. The SAP egress policy determines the QoS treatment to packets on service egress.

When the policy is created, by default only one queue is created. The user can create up to eight queues and associate each queue to different FCs on the SAPs to which this **sap-egress** policy is attached. A SAP egress policy allows the user to define the queue parameters for the eight queues.

Default

1

Parameters

policy-id

Specifies the ID of the SAP Egress policy.

Values 1 to 65535

create

Keyword to create a SAP Egress policy

remark

Syntax

remark *policy-id*

no remark

Context

config>qos>sap-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command specifies the remarking policy ID to use for marking packets on service egress (also known as SAP egress).

The remarking policy ID must be associated with the appropriate sap-egress policy and remarking must be enabled in the sap-egress policy to enable marking of packets sent out of the SAP. Only remarking policy of type dot1p, or dot1p-lsp-exp-shared is allowed to be used when the remark policy is associated with **sap-egress**. See [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information on remark policies.



Note:

On the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12, remarking is user configurable and supports SAP-based marking per SAP and port-based marking per access port. See [Access egress QoS policies on 7210 SAS-Mxp](#) and [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information about access port-based marking capabilities.

The **no** form of this command removes the explicit association of the remark policy and associates the default remark policy. If remarking is enabled and no remark policy is executed, the default remark policy is used to mark packets sent out. If no remark policy is executed and remarking is disabled, packets are not remarked.

Parameters

policy-id

Specifies the remark policy.

Values 1 to 65535

remarking

Syntax

no remarking

remarking

Context

config>qos>sap-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command enables the system to remark egress packets sent out of service egress (also known as Access SAP egress).

When remarking is enabled, the remark policy configured in the QoS policy context is used to determine the FC to QoS bit mapping. For example, when remarking is enabled in the -egress QoS policy, the remark policy associated with sap-egress QoS policy is used to determine the FC to dot1p mapping to use for marking packets sent out of access ports.

See [Remark policies](#) and [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information about the remark policy that can be used to configure FC to priority bit markings in different QoS policies associated with different service entities.

The **no** form of this command disables remarking.

Default

no remarking

scope

Syntax

scope {exclusive | template}

no scope

Context

config>qos>sap-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

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

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

Default

template

Parameters

exclusive

Specifies that the policy can only be applied to one interface or port. If a policy with an **exclusive** scope is assigned to a second interface, an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface. The system default policies cannot be put into the **exclusive** scope. An error will be generated if **scope exclusive** is executed in any policies with a *policy-id* equal to default QoS policies are configured with **template** scope. An error is generated if you try to modify the scope parameter from **template** to **exclusive** scope on default policies.

template

Specifies that the policy can be applied to multiple interfaces or ports on the router.

queue-mgmt

Syntax

[no] queue-mgmt name

Context

config>qos>sap-egress>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command associates the specified queue management policy with this queue.

The queue management policy is used to specify the queue buffer parameters and queue slope policy parameters.

The **no** form of this command associates the default SAP egress queue management policy with this queue.

Parameters

name

Specifies the name of the queue management policy, up to 32 characters.

queue-mode

Syntax

[no] queue-mode *queue-mode*

Context

config>qos>sap-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command determines whether the queue operates in strict or weighted mode.

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

Default

weighted

Parameters

queue-mode

Specifies the queue mode.

strict — Specifies that the scheduler schedules the queue in order of their priority in the two passes, the CIR loop and the PIR loop.

weighted — Specifies that the scheduler examines these queues in two passes - CIR loop and PIR loop. In the CIR loop, it distributes the available bandwidth to all the strict and then weighted queues in round-robin up to the configured CIR rate. It examines the weighted queues in the PIR loop, after examining all the strict queues and distributes the available bandwidth, if any, in the proportion of the configured weights.

rate

Syntax

rate cir *cir-rate-in-kbps* [**pir** *pir-rate-in-kbps*]

no rate

Context

config>qos>sap-egress>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the 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.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created on the access ports.

The **no** form of this command reverts all queues created with the queue ID by association with the QoS policy to the default PIR (max) and CIR (0) parameters.

Default

rate cir 0 pir max

Parameters

cir cir-rate-in-kbps

Specifies the administrative CIR rate, in kilobits, for the queue. The **cir** parameter overrides the default administrative CIR used by the queue. If the **rate** command is not executed or the **cir** parameter is not explicitly specified, the default CIR value is used.

Values 0 to 10000000, max

Default 0

pir pir-rate-in-kbps

Specifies the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a PIR setting is optional. If the **rate** command is not executed, the default PIR of maximum value is used. The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the *pir-rate* value.

Values 1 to 10000000, max

Default max

weight

Syntax

[no] weight *weight*

Context

config>qos>sap-egress>queue

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures the weight of the queue.

For queues configured in weighted mode, the CIR of the queues are met if bandwidth is available and the configured weights are considered in the PIR loop. The configured weight determines the proportion of available bandwidth that is given to this queue in comparison to other queues contending for bandwidth at the same priority level.

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

Default

1

Parameters

weight

Specifies the value for the **weight** parameter. The value is an integer value which specifies the proportion of available bandwidth to be allocated to this queue relative to other queues.

Values 1 to 15

scope

Syntax

scope {exclusive | template}

no scope

Context

config>qos>sap-egress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

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

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

Default

template

Parameters

exclusive

Specifies that the policy can only be applied to one interface port. If a policy with an **exclusive** scope is assigned to a second interface an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface. The system default policies cannot be put into the **exclusive** scope. An error will be generated if **scope exclusive** is executed in any policies with a policy ID equal to 1.

template

Specifies that the policy can be applied to multiple interface ports on the router.

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

12.2.2.2 Operational commands

copy

Syntax

copy sap-egress *src-pol dst-pol* [**overwrite**]

Context

config>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command copies the existing SAP egress QoS policy entries to another SAP egress 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 the destination policy already exists, the **overwrite** keyword must be specified.

Parameters

src-pol

Specifies the source policy.

Values 1 to 65535

dst-pol

Specifies the destination policy.

Values 1 to 65535

overwrite

Keyword to overwrite the information in the destination policy by the information in the source policy.

12.2.2.3 Show Commands

sap-egress

Syntax

sap-egress [*policy-id*] [association | detail]

Context

show>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays SAP egress QoS policy information.

Parameters

policy-id

Displays the policy ID of the sap-egress policy.

association

Displays associations related to the specified sap-egress policy.

detail

Displays detailed policy information including the policy associations.

Output

The following outputs are examples of SAP egress QoS policy information. The associated tables describe the output fields.

- [Sample output \(7210 SAS-Mxp\), Table 71: Output fields: SAP egress QoS policy \(7210 SAS-Mxp\)](#)
- [Sample output \(7210 SAS-R6 and 7210 SAS-R12\), Table 72: Output fields: SAP egress QoS policy \(7210 SAS-R6 and 7210 SAS-R12\)](#)

Sample output (7210 SAS-Mxp)

```
A:7210SAS>show>qos# sap-egress

=====
Sap Egress Policies
=====
-----
Policy-id Remark Remark Scope Name Description
```



```

Policy-id
-----
1          False 2          Template default          Default SAP egress
=====
A:7210SAS>show>qos#

```

Table 71: Output fields: SAP egress QoS policy (7210 SAS-Mxp)

Label	Description
Policy-Id	The ID that uniquely identifies the policy
Remark	True — Remarking is enabled for all the dot1q-tagged packets that egress the ports on which the sap-egress QoS policy is applied and remarking is enabled False — Remarking is disabled for the policy
Remark Pol Id	Displays the policy ID of the remarking policy
Accounting	Specifies whether the accounting mode is packet-based or frame-based
Scope	Exclusive — Implies that this policy can be applied only to a single access egress port Template — Implies that this policy can be applied to multiple access ports on the router
Description	A text string that helps identify the policy's context in the configuration file

Sample output (7210 SAS-R6 and 7210 SAS-R12)

```

A:SAS>config>qos# show qos sap-egress 1 detail

=====
QoS Sap Egress
=====
-----
Sap Egress Policy (1)
-----
Scope                : Template
Remark               : False          Remark Pol Id       : 2
Accounting           : frame-based
Description          : Default SAP egress QoS policy.
-----
Queue Rates and Rules
-----
-----
QueueId      CIR      CIR Adpt Rule      PIR      PIR Adpt Rule
-----
Queue1       0        closest           max       closest
Queue2       0        closest           max       closest
Queue3       0        closest           max       closest
Queue4       0        closest           max       closest
Queue5       0        closest           max       closest
Queue6       0        closest           max       closest
Queue7       0        closest           max       closest

```

Queue8	0	closest	max	closest

Parent Details				

QueueId	Port	CIR Level	PIR Weight	

Queue1	True	1	1	
Queue2	True	1	1	
Queue3	True	1	1	
Queue4	True	1	1	
Queue5	True	1	1	
Queue6	True	1	1	
Queue7	True	1	1	
Queue8	True	1	1	

High Slope				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	70	90	75
Queue2	Down	70	90	75
Queue3	Down	70	90	75
Queue4	Down	70	90	75
Queue5	Down	70	90	75
Queue6	Down	70	90	75
Queue7	Down	70	90	75
Queue8	Down	70	90	75

Low Slope				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	50	75	75
Queue2	Down	50	75	75
Queue3	Down	50	75	75
Queue4	Down	50	75	75
Queue5	Down	50	75	75
Queue6	Down	50	75	75
Queue7	Down	50	75	75
Queue8	Down	50	75	75

Burst Sizes and Time Average Factor				

QueueId	CBS	MBS	Time Average Factor	Queue-Mgmt

Queue1	def	def	7	default
Queue2	def	def	7	default
Queue3	def	def	7	default
Queue4	def	def	7	default
Queue5	def	def	7	default
Queue6	def	def	7	default
Queue7	def	def	7	default
Queue8	def	def	7	default

Associations				

Service-Id	: 1 (Epipe)		Customer-Id	: 1
- SAP	: 1/1/1:1			
Service-Id	: 101 (Epipe)		Customer-Id	: 1
- SAP	: 1/1/2:101			

```
-----
Mirror SAPs
-----
No Mirror SAPs Found.
=====
A:SAS>config>qos#
```

Table 72: Output fields: SAP egress QoS policy (7210 SAS-R6 and 7210 SAS-R12)

Label	Description
Policy-Id	Displays the ID that uniquely identifies the policy
Remark	True — Remarking is enabled for all the dot1q-tagged packets that egress the ports on which the SAP-egress QoS policy is applied and remarking is enabled False — Remarking is disabled for the policy
Remark Pol Id	Displays the policy ID of the remarking policy
Accounting	Specifies whether the accounting mode is packet-based or frame-based
Scope	Exclusive — Specifies that this policy can be applied only to a single access egress port Template — Specifies that this policy can be applied to multiple access ports on the router
Description	A text string that helps identify the policy context in the configuration file
Queue Rates and Rules	
QueueId	Displays the queue identifier associated with the SAP-egress QoS policy
Explicit/Default	Explicit — Specifies the egress dot1p bits marking for <i>fc-name</i> , if explicitly configured Default — Specifies the default dot1p value according to FC-Dot1p marking map as listed in Table 29: Default remarking policy for dot1p on 7210 SAS-R6 and 7210 SAS-R12 , if explicit values are not configured
CIR	Specifies the administrative 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 Adpt Rule	min — The operational CIR for the queue is equal to or greater than the administrative rate specified using the rate command max — The operational CIR for the queue is less than or equal to the administrative rate specified using the rate command

Label	Description
	closest — The operational CIR for the queue is the rate closest to the rate specified using the rate command without exceeding the operational PIR
PIR	Specifies the administrative PIR parameters for the queue The PIR defines the maximum rate that the queue can transmit packets through the access port.
PIR Adpt Rule	min — The operational PIR for the queue is equal to or greater than the administrative rate specified using the rate command max — The operational PIR for the queue is less than or equal to the administrative rate specified using the rate command closest — The operational PIR for the queue is the rate closest to the rate specified using the rate command
Parent Details	
QueueId	Displays the queue identifier associated with the SAP-egress QoS policy
Port	Indicates whether the parent scheduler is a port scheduler
CIR Level	Displays the priority of the queue in the CIR loop
PIR Weight	Displays the weight of the queue used in the PIR loop
High Slope/Low slope	
QueueId	Displays the queue identifier associated with the SAP-egress QoS policy
State	Displays the state of the queue The state of the queue can be either "Up" or "Down".
Start Avg	Specifies the low priority or high priority RED slope position for the shared buffer average utilization value, where the packet discard probability starts to increase above zero
Max Avg	Specifies the percentage of the shared buffer space for the buffer pool at which point the drop probability becomes "1" This parameter is expressed as a decimal integer.
Max Prob	Specifies the high-priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one
Burst Sizes and Time Average Factor	
QueueId	Displays the queue identifier associated with the SAP-egress QoS policy

Label	Description
CBS	Displays the configured CBS value
MBS	Displays the configured MBS value
Time Average Factor	Displays the value of the time average factor in use
Queue-Mgmt	Displays the queue management policy in use
Service Associations	
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 SAP within the service where the SAP-egress policy is applied

13 QoS port scheduler policies for 7210 SAS-T

This section provides information to configure port scheduler policies using the command line interface.

13.1 Overview

This section provides an overview of QoS port scheduler policies on the 7210 SAS-T.

13.1.1 Configuring port scheduler policies

The **port-scheduler-policy** command creates a port scheduler template which may be assigned to an egress port. Only one port scheduler policy is allowed per port. There is a "default" port-scheduler policy (which services the queues of the port in a Strict order) associated with each port. To change the behavior, users can associate the port with another port-scheduler policy. The policy contains mode commands to set the mode of scheduling (RR, Strict, WRR, WDRR) and queue commands to set the weight of the queue, when the mode is set to one of WRR and WDRR. In WRR/WDRR, a **strict** option treats that particular queue as a strict queue, this leads to a hybrid mode of scheduling (WRR+Strict, WDRR+Strict).

13.1.2 Basic configurations

A basic QoS port scheduler policy must conform to the following:

- Each QoS port scheduler policy must have a unique policy name.

13.1.2.1 Creating a QoS port scheduler policy

To create a port scheduler policy, define the following:

- a port scheduler policy name
- include a description - provides a brief overview of policy features

Use the following syntax to create a QoS port scheduler policy.

Note that the **create** keyword is included in the command syntax upon creation of a policy.

```
config>qos
  port-scheduler-policy port-scheduler-name [create]
    description description-string
    mode {strict | rr | wrr | wdr}
    queue queue-id [strict | weight weight]
```

Output example

The following is a sample port scheduler policy configuration output.

```
*A:card-1>config>qos>port-sched-plcy# info
```

```

-----
mode WRR
queue 1 weight 1
queue 2 weight 3
queue 3 weight 5
queue 5 weight 5
queue 6 weight 1
-----
*A:card-1>config>qos>port-sched-plcy#

```

13.2 Service management tasks

This section describes the service management tasks.

13.2.1 Copying and overwriting scheduler policies

You can copy an existing QoS policy, rename it with a new QoS policy value, or overwrite an existing policy. The overwrite option must be specified or an error occurs if the destination policy exists.

```
config>qos> copy port-scheduler-policy src-name dst-name [overwrite]
```

Output example

```

*A:Dut-1>config>qos# port-scheduler-policy psp create
*A:Dut-1>config>qos>port-sched-plcy# mode wdr
*A:Dut-1>config>qos>port-sched-plcy# queue 1 weight 1
*A:Dut-1>config>qos>port-sched-plcy# queue 2 weight 2
*A:Dut-1>config>qos>port-sched-plcy# queue 3 weight 5
*A:Dut-1>config>qos>port-sched-plcy# info
-----
mode wdr
queue 2 weight 2
queue 3 weight 5
-----
*A:Dut-1>config>qos>port-sched-plcy# exit
*A:Dut-1>config>qos# exit
*A:Dut-1>config# qos copy port-scheduler-policy psp psp1
*A:Dut-1>config# qos copy port-scheduler-policy psp psp1
MINOR: CLI Destination "psp1" exists - use {overwrite}.

*A:Dut-1>config# show qos port-scheduler-policy
=====
Port Scheduler Policies
=====
Policy-Id          Description          Mode
-----
default           Default Port Scheduler policy.  STRICT
psp                WDRR
psp1               WDRR
=====
*A:Dut-1>config#

*A:Dut-1>config# show qos port-scheduler-policy psp
=====
QoS Port Scheduler Policy

```

```
=====
Policy-Name      : psp
Accounting       : packet-based
Mode            : WDRR
Last changed    : 04/12/2001 02:04:16
Queue 1 Weight: : 1
Queue 2 Weight: : 2
Queue 3 Weight: : 5
Queue 4 Weight: : 1
Queue 5 Weight: : 1
Queue 6 Weight: : 1
Queue 7 Weight: : 1
Queue 8 Weight: : 1
=====

*A:Dut-1>config#

*A:Dut-1>config# show qos port-scheduler-policy psp1
=====
QoS Port Scheduler Policy
=====
Policy-Name      : psp1
Accounting       : packet-based
Mode            : WDRR
Last changed    : 04/12/2001 02:05:00
Queue 1 Weight: : 1
Queue 2 Weight: : 2
Queue 3 Weight: : 5
Queue 4 Weight: : 1
Queue 5 Weight: : 1
Queue 6 Weight: : 1
Queue 7 Weight: : 1
Queue 8 Weight: : 1
=====

*A:Dut-1>config#
```

13.2.2 Editing QoS policies

To prevent configuration errors use the copy command to make a duplicate of the original policy to a work area, make the edits, and then overwrite the original policy.

14 Schedulers on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE

This section describes the port scheduler and its behavior for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE. The port scheduler on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE is configured using the port scheduler policies.

14.1 Configuring scheduler policies

On 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, there are 2 levels of egress schedulers – a port egress scheduler operating at line-rate or the configured port egress rate and a per FC egress scheduler which provides an option to the user to configure CIR and PIR rate. The FC egress scheduler is parented by the port egress scheduler. A port has a single port egress scheduler with a scheduling node for each FC, for a total of eight (8), as its children. Each FC egress scheduling node has 2 queues as its children, one for unicast traffic and one for multicast traffic (all BUM traffic):

- The rate parameter configured under a queue in the access-egress policy and network queue policy limits the amount of bandwidth that the FC will get. The queue ID in the policy maps one-to-one to an FC (with queue #8 mapping to FC nc, queue #7 to FC h2, and so on to queue #1 to FC be). The amount of bandwidth given to a FC is equally distributed among the unicast queue and multicast queue in a WDRR fashion (with equal weights assigned to the queues).

The port scheduler policy mode parameter determines the scheduling order of the FCs as follows:

- **Mode = Strict (aka strict priority scheduling)**

The port egress scheduler operates in two loops – CIR loop and PIR loop. If configured in strict priority, FC nc is scheduled before FC 'be' in both CIR loop and PIR loop. Lower priority FC, will be allocated bandwidth only if the higher priority FC has no traffic to send in the given loop. In other words, in the CIR loop only after meeting the CIR rate of the higher priority FC, the scheduler schedules the lower priority FC, only if port bandwidth remains. After completing the CIR loop and satisfying the CIR rates configured for all the FCs, the scheduler starts the PIR loop (only if more port bandwidth remains). In the PIR loop too, bandwidth is first allocated to higher priority FC followed by lower priority FCs. It is recommended to configure the rates correctly for higher priority FC when using strict priority mode to avoid starvation of lower priority FC.

- **Mode – weighted (WRR or WDRR)**

If configured for weighted mode (either WDRR or WRR), the weight configured determines the amount of bandwidth a FC gets when contending for traffic with other FCs in the PIR loop. The port scheduler operates in two loop – CIR loop and PIR loop. In CIR loop, the available bandwidth is distributed in round-robin order to all the FCs whose CIR rate is not met. Weight is not considered in the CIR loop. In the PIR loop, the port scheduler distributes the available bandwidth to all the FCs in proportion to the configured weight until their PIR is met. The accounting for amount of bandwidth distributed is in terms of packets which results in unfair advantage to FCs that receive more number of bigger packets.

- **Mode – Round-Robin (RR)**

If configured for round-robin mode, the port scheduler operates in two loop – CIR loop and PIR loop. In CIR loop, the available bandwidth is distributed in round-robin fashion to all the FCs whose CIR rate

is not met. Once the CIR is met for all FCs, in the PIR loop, the port scheduler distributes the available bandwidth to all the FCs until their PIR is met. The accounting for amount of bandwidth distributed is in terms of packets which results in unfair advantage to FCs that receive only larger size packets.

- **Mode – Hybrid (Strict + WDRR or WRR)**

If the WRR/WDRR weight associated with a particular FC is set to strict, the FC is considered to be operating in a strict priority mode. The set of strict priority queues is serviced first in the order of their priority, with higher priority FC not scheduled before FC be. In this mode, the scheduler services the strict FCs, followed by the FCs configured with weights in both the CIR and PIR loop. The scheduler ensures that it meets the CIR of all the FCs (both FCs configured as strict and FCs configured as weighted), before scheduling the FCs in the PIR loop (assuming sufficient port bandwidth is available). If multiple FCs are configured as strict, the higher-priority strict queues are serviced first before the lower priority strict queues in both the CIR and the PIR loop. The weights configured for the FCs are only considered during the PIR loop to distribute the available bandwidth in proportion to the weights. Care must be taken when configuring strict priority queues to avoid starvation of lower priority strict queues or weighted queues.

14.2 QoS port scheduler policy command reference

14.2.1 Command hierarchies

- [Port scheduler policy configuration commands](#)
- [Operational commands](#)
- [Show commands](#)

14.2.1.1 Port scheduler policy configuration commands

```
config
- qos
  - [no] port-scheduler-policy port-scheduler-name [create]
    - description description-string
    - no description
    - mode {strict | rr | wrr | wdr}
    - no mode
    - queue queue-id [strict | weight weight]
    - no queue queue-id
```

14.2.1.2 Operational commands

```
config
- qos
  - copy port-scheduler-policy src-name dst-name [overwrite]
```

14.2.1.3 Show commands

```
show
- qos
- port-scheduler-policy [port-scheduler-policy-name] [association]
```

14.2.2 Command descriptions

14.2.2.1 Configuration commands

14.2.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>port-scheduler-policy

Platforms

7210 SAS-T (in access-uplink mode and network mode), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

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

description-string

Specifies 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.

14.2.2.2 Port scheduler policy commands

port-scheduler-policy

Syntax

[no] port-scheduler-policy *port-scheduler-name* [**create**]

Context

config>qos

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

The default scheduling done for a port is strict scheduling. When a port-scheduler policy is applied to a port, it overrides the default scheduling and determines the type of scheduling (Strict, RR, WRR, WDRR, WRR/WD RR + Strict) to be done between the 8 CoS queues of that particular port. When a port scheduler policy is detached from a port, the port reverts back to the default scheduling (strict).

The **no** form of this command removes the policy from the system.

Parameters

port-scheduler-name

Specifies an existing policy name. Each port-scheduler policy name should be unique and can go up to 32 ASCII characters in length.

create

Keyword to create a port scheduler policy.

mode

Syntax

mode {**strict** | **rr** | **wrr** | **wdr**}

no mode

Context

config>qos>port-sched-plcy

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command configures a particular mode of scheduling for the policy. For example, this implies that when a policy with a mode RR is applied to a port then that port will follow the round robin type of scheduling between its queues.

Parameters

mode

Specifies the port scheduler policy mode.

strict — Strict scheduler mode

rr — Round Robin

wrr — Weighted Round Robin

wdrr — Weighted Deficit Round Robin

queue

Syntax

queue *queue-id* [**strict** | **weight** *weight*]

no queue *queue-id*

Context

config>qos>port-sched-plcy

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command configures a port scheduler queue. The queue and its weights can be configured only for WRR/WDRR modes. The weight specified in case of WRR corresponds to the number of packets that needs to be sent out in a cycle for that particular queue.

For WDRR, the weight specified is the ratio of traffic that will be sent out for that particular queue. For example, in WDRR, if a weight value for queue 1 is 1 and a weight value for queue 2 is 5, then traffic out of the port is in the ratio of 1:5 between the queues (1 and 2) provided no traffic is flowing in the other queues. If the keyword **strict** is specified in any of the queues, then that particular queue will be treated as strict. This set of strict priority queues is serviced first in the order of their CoS numbering (the higher numbered CoS queue receives service before smaller numbered queues).



Note:

On the 7210 SAS-Sx 1/10GE: standalone and standalone-VC, the queue ID represents the FC. The FC is determined by the FC-to-queue ID map. For more information, see [Schedulers on 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE](#).

The **no** form of this command under a WRR/WDRR mode will set the queue weights to default; for example, 1.

Parameters

queue-id

Specifies the queue ID.

Values 1 to 8

strict

Specifies strict access.

weight *weight*

Specifies the number of packets in case of WRR and ratio of traffic out in WDRR.

Values 1 to 15

14.2.2.3 Operational commands

copy

Syntax

copy port-scheduler-policy *src-name dst-name* [overwrite]

Context

config>qos

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command copies existing port scheduler QoS policy entries for a port scheduler QoS policy to another port scheduler 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

port-scheduler-policy src-name dst-name

Specifies the source policy that the **copy** command will attempt to copy from and specifies the destination policy name to which the command will copy a duplicate of the policy. This parameter indicates that the source policy and the destination policy are port scheduler policy IDs.

overwrite

Keyword to replace the destination policy name. When the **overwrite** keyword is specifies, everything in the existing destination policy will be completely overwritten with the contents of the source policy.

14.2.2.4 Show commands

port-scheduler-policy

Syntax

port-scheduler-policy [*port-scheduler-policy-name*] [**association**]

Context

show>qos

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command displays port-scheduler policy information

Parameters

port-scheduler-policy-name

Displays information for the specified existing port scheduler policy.

association

Displays associations related to the specified port scheduler policy.

Output

The following outputs are examples of QoS port scheduler policy information, and the associated tables describe the output fields:

- [Sample output \(7210 SAS-T\)](#), [Table 73: Output fields: port scheduler policy for 7210 SAS-T](#)
- [Sample output \(7210 SAS-Sx 1/10GE\)](#), [Table 74: Output fields: port scheduler policy for 7210 SAS-Sx 1/10GE](#)
- [Sample output \(association\)](#), [Table 75: Output fields: association](#)

Sample output (7210 SAS-T)

*A:Dut-1>config# show qos port-scheduler-policy		
=====		
Port Scheduler Policies		
=====		
Policy-Id	Description	Mode

default	Default Port Scheduler policy.	STRICT

```

psp
psp1
=====
WDRR
WDRR
=====
*A:Dut-1>config#

*A:Dut-1>config# show qos port-scheduler-policy psp association
=====
QoS Port Scheduler Policy
=====
Policy-Name      : psp
Accounting       : packet-based
Mode             : WDRR

-----
Associations
-----
- Port : 1/1/1

=====
*A:Dut-1>config#

*A:Dut-1>config# show qos port-scheduler-policy psp
=====
QoS Port Scheduler Policy
=====
Policy-Name      : psp
Accounting       : packet-based
Mode             : WDRR
Last changed     : 04/12/2001 02:04:16
Queue 1 Weight:  : 1
Queue 2 Weight:  : 2
Queue 3 Weight:  : 5
Queue 4 Weight:  : 1
Queue 5 Weight:  : 1
Queue 6 Weight:  : 1
Queue 7 Weight:  : 1
Queue 8 Weight:  : 1
=====
*A:Dut-1>config#

```

Table 73: Output fields: port scheduler policy for 7210 SAS-T

Label	Description
Policy-Id	The ID of the policy
Description	Description of the policy
Mode	Displays the port scheduler policy mode (STRICT, RR, WRR, WDRR)
Associations	Displays associations related to the specified port scheduler policy
Policy-Name	Displays the port scheduler policy name
Accounting	Displays whether the accounting mode is frame-based or packet-based

Label	Description
Last Changed	Displays the last time the configuration changed
Queue #	Displays the weight of the queue, if configured

Sample output (7210 SAS-Sx 1/10GE)

```
*A:K-SAS-Sx>config>qos# show qos port-scheduler-policy
=====
Port Scheduler Policies
=====
Policy-Id          Description                                     Mode
-----
1                  Default Port Scheduler policy.                STRICT
default            Default Port Scheduler policy.                STRICT
=====
*A:K-SAS-Sx>config>qos# show qos port-scheduler-policy "default" association
=====
QoS Port Scheduler Policy
=====
Policy-Name       : default
Description       : Default Port Scheduler policy.
Accounting        : packet-based
Mode              : STRICT

-----
Associations
-----
- Port : 1/1/1
- Port : 1/1/2
- Port : 1/1/3
- Port : 1/1/4
- Port : 1/1/5
- Port : 1/1/6
- Port : 1/1/7
.....
=====
*A:K-SAS-Sx>config>qos# show qos port-scheduler-policy "default"
=====
QoS Port Scheduler Policy
=====
Policy-Name       : default
Description       : Default Port Scheduler policy.
Accounting        : packet-based
Mode              : STRICT
Last changed     : 01/03/2000 07:34:35
Number Of Queues  : 8
=====
```

Table 74: Output fields: port scheduler policy for 7210 SAS-Sx 1/10GE

Label	Description
Policy-Id	The ID of the policy
Description	Description of the policy

Label	Description
Mode	Displays the port scheduler policy mode (STRICT, RR, WRR, WDRR)
Associations	Displays associations related to the specified port scheduler policy
Policy-Name	Displays the port scheduler policy name
Accounting	Displays whether the accounting mode is frame-based or packet-based
Last Changed	Displays the last time the configuration changed
Number of Queues	Displays the number of queues, if configured

Sample output (association)

```
*A:card-1# show qos port-scheduler-policy default association
=====
QoS Port Scheduler Policy
=====
Policy-Name      : default
Description      : Default Port Scheduler policy.
Accounting       : packet-based
Mode             : STRICT

-----
Associations
-----
- Port : 1/1/3
- Port : 1/1/6
- Port : 1/1/7
- Port : 1/1/8
- Port : 1/1/9
- Port : 1/1/10
- Port : 1/1/11
- Port : 1/1/12
- Port : 1/1/13
- Port : 1/1/14
- Port : 1/1/16
- Port : 1/1/17
- Port : 1/1/18
- Port : 1/1/19
- Port : 1/1/21
- Port : 1/1/22
- Port : 1/1/23
- Port : 1/1/24
...
=====
*A:card-1#

*A:Dut-1>config# show qos port-scheduler-policy default
=====
QoS Port Scheduler Policy
=====
Policy-Name      : default
Description      : Default Port Scheduler policy.
Accounting       : packet-based
```

```
Mode          : STRICT
Last changed  : 04/11/2001 19:59:21
Number Of Queues : 8
```

```
=====
*A:Dut-1>config#
```

Table 75: Output fields: association

Label	Description
Policy-Name	Displays the port scheduler policy name
Description	Description of the policy
Accounting	Displays whether the accounting mode is frame-based or packet-based
Mode	Displays the port scheduler policy mode (STRICT, RR, WRR, WDRR)
Associations	Displays associations related to the specified port scheduler policy
Last Changed	Displays the last time the configuration changed
Number of Queues	Displays the number of queues, if configured

15 Schedulers on 7210 SAS-Mxp

This section provides information about the scheduler support available in the 7210 SAS-Mxp devices for network port and SAP.

15.1 Overview

On the 7210 SAS-Mxp, users have an option to use either port-based egress queuing and shaping or SAP-based egress queuing and shaping for SAPs configured on access ports or hybrid ports. The **configure system resource-profile qos port-scheduler-mode** command allows you to select the mode to be used for SAPs configured on all the ports of the node (that is, this is a per node setting). The following sections describe the behavior of the scheduler in these two modes.



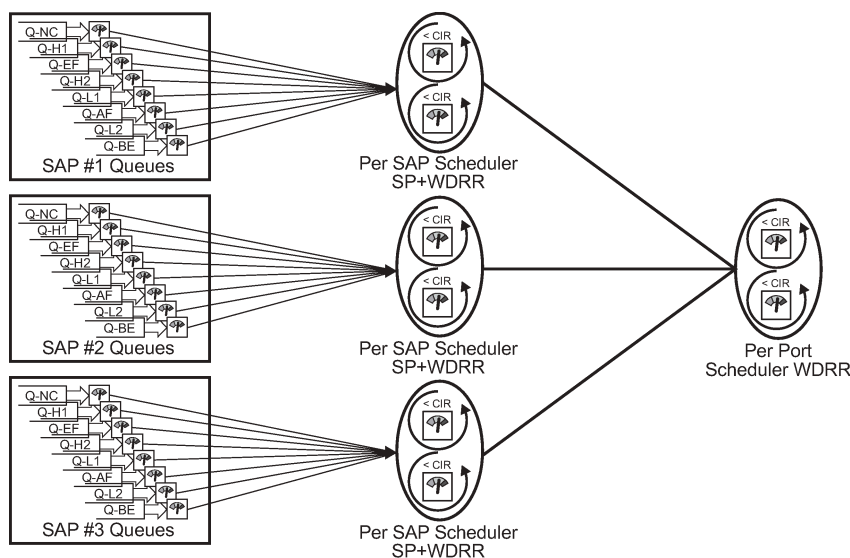
Note:

The queue parameters and scheduling parameters for the queue are configured in the SAP egress policies when using SAP-based egress queues, and are configured in the access-egress policies when using port-based egress queues.

15.2 Scheduling with SAP-based queues on access ports

When SAP-based scheduling is enabled, the following figure shows the scheduling for access port, with multiple SAPs configured and when the **port-scheduler-mode** is disabled. There are 8 egress queues per SAP, a per SAP scheduler and a per access port scheduler.

Figure 6: SAP egress scheduling



hw0548



Note:

- Each FC/queue for the port can be shaped to configured rates (CIR/PIR). This is used to control the amount of bandwidth allocated to the FC/queue.
- FC-to-queue mapping is system-defined and not user configurable.
- The queue number determines the priority of the queue, which is used only when the queues are configured as "strict." Queue "8" is the highest priority, and queue "1" is the lowest priority.
- On the 7210 SAS-Mxp, only unicast traffic sent out of RVPLS SAPs uses per-SAP egress queues. BUM traffic sent out of RVPLS SAPs uses per-port egress queues. Per-port egress queues are not depicted in the preceding diagram. There are eight per-port queues, and they contend with per-SAP queues for bandwidth.
- A queue can be defined to operate in **strict** mode or **weighted** mode. The queue mode determines the order of scheduling by the port scheduler.
- In a CIR loop, scheduling is packet-based round-robin with a weight of 1.

The behavior of the scheduler for an access port is as follows:


- Per port scheduler is available and works at line-rate or configured port egress rate.
- Port scheduler distributes the bandwidth available to all the SAPs using WDRR scheduling mechanism (that is, all SAPs have equal weights assigned by the system).
- The port scheduler uses the following two passes to distribute bandwidth across SAPs.
 - **CIR loop**
The port scheduler distributes the available bandwidth to all the SAP in a round-robin order up to the configured CIR rate (CIR is configured in the aggregate shaper rates for the SAP).
 - **PIR loop**
The port scheduler distributes the remaining bandwidth (the bandwidth available after the CIR loop) to all SAPs in a round-robin order (all SAPs are assigned equal weights by the system) up to the configured PIR rate (PIR is configured in the aggregate shaper rates for the SAP).
- Each SAP has a per SAP scheduler which operates in SP + WDRR mode and an aggregate per SAP shaper (CIR/PIR). The per SAP scheduler distributes the available bandwidth to the configured strict and weighted SAP queues, using the configured mode and rates, in 2 passes - CIR loop and PIR loop
- The CIR loop distributes the available bandwidth (from the bandwidth allocated to it by the port scheduler) to all the queues in the following order:
 - higher priority **strict** queues get the bandwidth up to the configured CIR
 - any remaining bandwidth, if available, is distributed among the lower priority **strict** queues up to the configured CIR
 - any remaining bandwidth, if available, is distributed among the **weighted** queues (in the CIR loop, weights are not used and therefore the bandwidth is distributed in equal proportion irrespective of weights configured)
- The PIR loop distributes the remaining bandwidth (the bandwidth remaining after CIR loop) to all the queues in the following order:
 - higher priority **strict** queues get the bandwidth, up to the configured PIR
 - any remaining bandwidth, if available, is distributed among the lower priority **strict** queues, up to the configured PIR

- any remaining bandwidth, if available, is distributed among the **weighted** queues in proportion to their configured weights
- Each queue can be configured with a queue-mode (**strict** or **weighted**) and is associated with a shaper (which allows for configuration of CIR/PIR). The queue mode determines the order of scheduling by the SAP scheduler and shaper rate controls the amount of bandwidth used by the queue.

15.3 Scheduling on network ports

For a network port, the scheduling behavior is similar, except that per SAP scheduler is not present in the hierarchy. Instead, per port scheduler distributes the available port bandwidth to all the queues configured on the port in two passes with the behavior being similar to the per SAP scheduler (as mentioned above). Additionally, all the traffic sent out on the network port uses a set of 8 queues which are mapped to the 8 forwarding classes. This is shown in the following figure.


Figure 7: Scheduling on network ports



15.4 Scheduling on hybrid port with SAP-based egress queues

The following figure shows the scheduling hierarchy for hybrid port with both SAPs and network interfaces configured.

Figure 8: Hybrid port egress scheduling



The scheduling behavior is very similar to the SAP-based queues on the access ports. That is, the network port queues used for network traffic are treated as another SAP node in the scheduling hierarchy.

15.4.1 Port-based scheduling and queuing on access ports

Figure 9: Port-based scheduling and queuing shows port-based scheduling and queuing enabled on Access-ports on 7210 SAS-Mxp.


When **port-scheduler-mode** is enabled, traffic sent out of SAPs configured on access ports and hybrid ports, share a set of 8 egress queues which are mapped to the 8 forwarding classes. A per port scheduler (similar to the one available for network port) distributes the available port bandwidth to all the queues configured on the port in two passes with the behavior being similar to the per SAP scheduler (see below for more details).



Note:

Per-SAP scheduler is not present in the scheduler hierarchy when **port-scheduler-mode** is enabled.

Figure 9: Port-based scheduling and queuing



The behavior of the Port-based Queuing and Scheduling on Access Ports is as follows:

- If enabled, all SAPs on the node/chassis use port-based queuing.

- In other words, user has an option to use either SAP-based queue for all SAPs configured on the node or port-based queues for all SAPs configured on the node. A mix and match of some SAPs using port-based queues and some SAPs using SAP-based queues is not supported.
- All SAPs on an access port share the 8 egress queues on the port. On Hybrid ports, SAPs use network port queues. In other words, on hybrid port, all the SAPs configured on the port and the network port IP interfaces share the 8 egress queues on the port.
- Supports 2 level hierarchical shaping, with Per queue shaper and per port aggregate shaper (ERL).
- Each FC/queue of the port can be shaped to configured rates (CIR/PIR). This is used to control the amount of bandwidth allocated to the FC/queue.
- FC to queue mapping is system-defined and not user configurable.
- The queue number determines the priority of the queue. Priority of the queue is used only when the queues are configured as **strict**. Queue "8" is the highest priority and Queue "1" is the lowest priority.
- A queue can be defined to operate in **strict** mode or **weighted** mode. The queue mode determines the order of scheduling by the port scheduler.

The scheduling behavior is similar to the one supported on SAPs (modified as below):


Each Access Port has a per port scheduler which operates in SP + WDRR mode and an aggregate per port shaper (ERL). The per port scheduler distributes the available bandwidth to the configured **strict** and **weighted** queues, using the configured mode and rates, in 2 passes - CIR loop and PIR loop:

- The CIR loop distributes the available bandwidth to all the queues in the following order:
 - Higher priority **strict** queues get the bandwidth up to the configured CIR.
 - Any remaining bandwidth, if available, is distributed among the lower priority **strict** queues up to the configured CIR.
 - Any remaining bandwidth, if available, is distributed among the **weighted** queues (in the CIR loop, weights are not used and therefore the bandwidth is distributed in equal proportion irrespective of weights configured).
- The PIR loop distributes the remaining bandwidth (the bandwidth remaining after CIR loop) to all the queues in the following order:
 - Higher priority **strict** queues get the bandwidth, up to the configured PIR.
 - Any remaining bandwidth, if available, is distributed among the lower priority **strict** queues, up to the configured PIR.
 - Any remaining bandwidth, if available, is distributed among the **weighted** queues in proportion to their configured weights.

15.5 Scheduling on hybrid port with port-based SAP queues

The following figure shows the scheduling hierarchy for a hybrid port with SAPs using port-based queues and network IP interfaces using port-based queues. The scheduling behavior is similar to that of access port when port-based queues are used (as described above).

Figure 10: Scheduling hierarchy for a hybrid port with SAPs



16 Schedulers on 7210 SAS-R6 and 7210 SAS-R12

This chapter provides information about the scheduler support available on the 7210 SAS-R6 and 7210 SAS-R12 devices for network ports and SAPs.



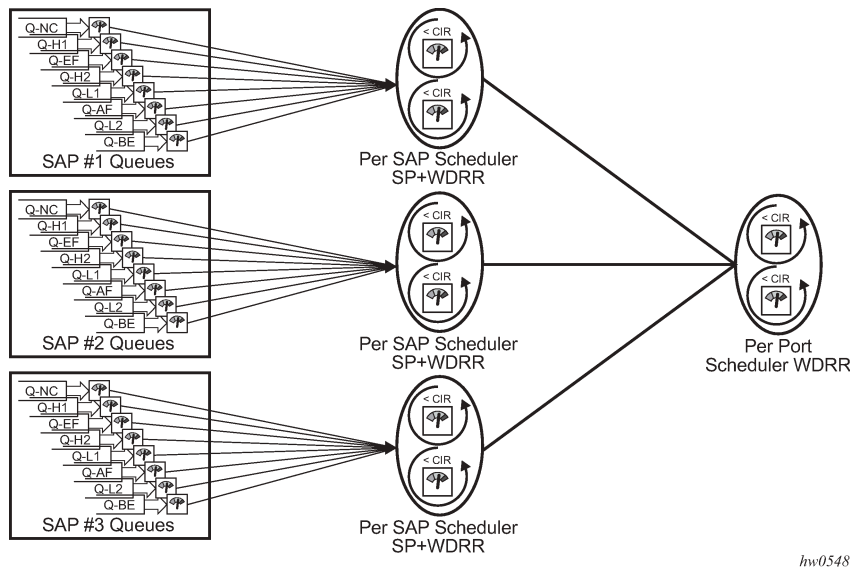
Note:

See [Service ingress QoS policies](#), [Service egress policies on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#), and [Network queue QoS policies](#) for more information about the configuration examples for scheduling parameters.

16.1 Scheduling with SAP-based queues on access ports

The following figure shows the scheduling for access ports when SAP-based scheduling is enabled, with multiple SAPs configured and when the port-scheduler-mode is disabled. There are eight egress queues per SAP, a per-SAP scheduler, and a per-access-port scheduler.

Figure 11: SAP egress scheduling



hw0548



Note:

- Each FC or queue of the port can be shaped to configured rates (CIR and PIR). This is used to control the amount of bandwidth allocated to the FC or queue.
- FC-to-queue mapping is system-defined and not user configurable.
- The queue number determines the priority of the queue. Priority of the queue is used only when the queues are configured as **strict**. Queue 8 is the highest priority and queue 1 is the lowest priority.

- On the 7210 SAS-R6 and 7210 SAS-R12, only unicast traffic sent out of RVPLS SAPs uses per-SAP egress queues. BUM traffic sent out of RVPLS SAPs uses per-port egress queues. These per-port egress queues are not shown in [Figure 11: SAP egress scheduling](#). There are eight per-port queues, which compete with per-SAP queues for bandwidth.
- A queue can be defined to operate in strict mode or weighted mode. The queue mode determines the order of scheduling by the port scheduler.
- In a CIR loop, scheduling is packet-based round-robin with a weight of 1.


The scheduler for an access port has the following behavior:

- The per-port scheduler is available and works at line-rate or configured port egress rate.
- The port scheduler distributes the bandwidth available to all the SAPs using WDRR scheduling mechanism (that is, all SAPs have equal weights assigned by the system).
- The port scheduler uses the following two passes to distribute bandwidth across SAPs:
 - **CIR loop**
The port scheduler distributes the available bandwidth to all the SAPs in round-robin order, up to the configured CIR rate (CIR is configured in the aggregate shaper rates for the SAP).
 - **PIR loop**
The port scheduler distributes the remaining bandwidth (the bandwidth available after the CIR loop) to all the SAPs in round-robin order (all SAPs are assigned equal weights by the system), up to the configured PIR rate (PIR is configured in the aggregate shaper rates for the SAP).
- Each SAP has a per-SAP scheduler that operates in SP + WDRR mode and an aggregate per-SAP shaper (CIR/PIR). The per-SAP scheduler distributes the available bandwidth to the configured strict and weighted SAP queues, using the configured mode and rates, in two passes: CIR loop and PIR loop.
- The CIR loop distributes the available bandwidth (from the bandwidth allocated to it by the port scheduler) to all the queues in the following order:
 - higher priority strict queues receive bandwidth up to the configured CIR
 - any remaining bandwidth, if available, distributed among the lower priority strict queues, up to the configured CIR
 - any remaining bandwidth, if available, distributed among the weighted queues (in the CIR loop, weights are not used and therefore the bandwidth is distributed in equal proportion irrespective of weights configured)
- The PIR loop distributes the remaining bandwidth (the bandwidth remaining after CIR loop) to all the queues in the following order:
 - higher priority strict queues receive the bandwidth, up to the configured PIR
 - any remaining bandwidth, if available, distributed among the lower priority strict queues, up to the configured PIR
 - any remaining bandwidth, if available, distributed among the weighted queues in proportion to their configured weights
- Each queue can be configured with a queue-mode (strict or weighted) and is associated with a shaper (which allows users to configure CIR and PIR). The queue mode determines the order of scheduling by the SAP scheduler and shaper rate controls the amount of bandwidth used by the queue.

16.2 Scheduling on network ports

For a network port, the scheduling behavior is similar to [Scheduling with SAP-based queues on access ports](#), except that a per-SAP scheduler is not present in the hierarchy. Instead, the per-port scheduler distributes the available port bandwidth to all the queues configured on the port in two passes using similar behavior to the per-SAP scheduler. Additionally, all traffic sent out on the network port uses a set of eight queues that are mapped to the eight forwarding classes. The following figure shows this behavior.

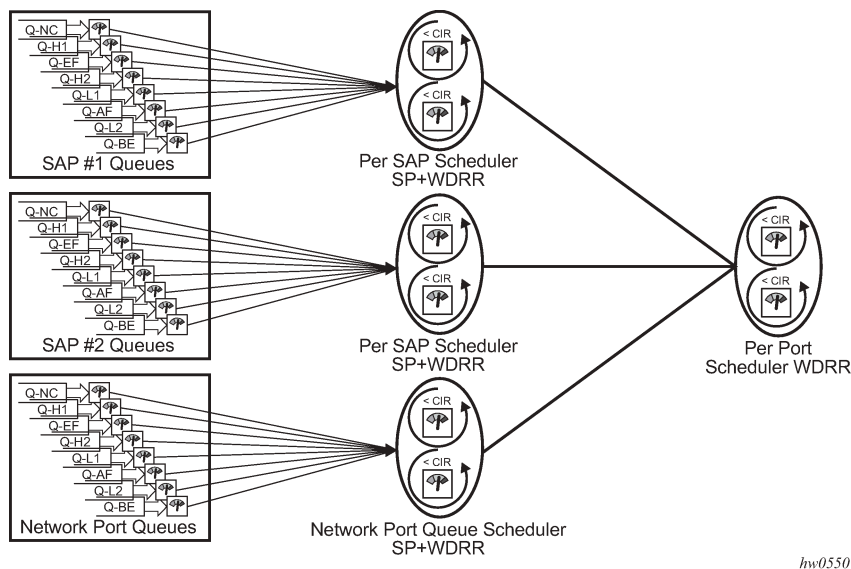
Figure 12: Scheduling on network ports



16.3 Scheduling on hybrid port with SAP-based egress queues

The following figure shows the scheduling hierarchy for hybrid ports with both SAPs and network interfaces configured.

Figure 13: Hybrid port egress scheduling



The scheduling behavior is similar to the SAP-based queues on the access ports. The network port queues used for network traffic are treated like a SAP node in the scheduling hierarchy.

16.3.1 Port-based scheduling and queuing on access ports

When port-scheduler mode is enabled, traffic sent out of SAPs configured on access ports and hybrid ports, share a set of eight egress queues that are mapped to the eight forwarding classes. A per-port scheduler (similar to the one available for network port) distributes the available port bandwidth to all the queues configured on the port in two passes using similar behavior to the per-SAP scheduler.




Note:

The per-SAP scheduler is not present in the scheduler hierarchy when port-scheduler mode is enabled.

The following figure shows the port-based scheduling and queuing enabled on access ports on the 7210 SAS-R6 and 7210 SAS-R12.

Figure 14: Port-based scheduling and queuing



Port-based queuing and scheduling on access ports has the following behavior:

- If enabled, all SAPs on the node use port-based queuing. This means that the user has an option to use either SAP-based queues for all SAPs configured on the node, or port-based queues for all SAPs configured on the node. A mix and match of some SAPs using port-based queues and some SAPs using SAP-based queues is not supported
- All SAPs on an access port share the eight egress queues on the port. On hybrid ports, SAPs use network port queues, which means that, on hybrid ports, all the SAPs configured on the port and the network port IP interfaces share the eight egress queues on the port.
- This functionality supports two-level hierarchical shaping, with per-queue shaper and per-port aggregate shaper (ERL).
- Each FC or queue of the port can be shaped to configured rates (CIR/PIR). This is used to control the amount of bandwidth allocated to the FC or queue.
- FC-to-queue mapping is system-defined and not user configurable.
- The queue number determines the priority of the queue. Priority of the queue is used only when the queues are configured as **strict**. Queue 8 is the highest priority and queue 1 is the lowest priority.
- A queue can be defined to operate in **strict** mode or **weighted** mode. The queue mode determines the order of scheduling by the port scheduler.

The scheduling behavior is similar to the scheduling behavior that is supported on SAPs with the following modifications:


- Each access port has a per-port scheduler that operates in SP + WDRR mode and an aggregate per-port shaper (ERL). The per-port scheduler distributes the available bandwidth to the configured **strict** and **weighted** queues, using the configured mode and rates, in two passes: CIR loop and PIR loop.
- The CIR loop distributes the available bandwidth to all the queues in the following order:

- Higher priority **strict** queues receive bandwidth up to the configured CIR.
- Any remaining bandwidth, if available, is distributed among the lower priority **strict** queues up to the configured CIR.
- Any remaining bandwidth, if available, is distributed among the **weighted** queues (in the CIR loop, weights are not used and therefore the bandwidth is distributed in equal proportion irrespective of weights configured)
- The PIR loop distributes the bandwidth remaining after CIR loop to all the queues in the following order:
 - Higher priority **strict** queues receive the bandwidth, up to the configured PIR.
 - Any remaining bandwidth, if available, is distributed among the lower priority **strict** queues, up to the configured PIR.
 - Any remaining bandwidth, if available, is distributed among the **weighted** queues in proportion to their configured weights.

16.4 Scheduling on hybrid port with port-based SAP queues

The following figure shows the scheduling hierarchy for the hybrid port, with SAPs using port-based queues and network IP interfaces using port-based queues. The scheduling behavior is similar to that of an access port when port-based queues are used (as described previously).

Figure 15: Scheduling hierarchy for the hybrid port



17 Slope QoS policies

This chapter provides information to configure slope QoS policies using the command line interface.

This chapter is applicable to the 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC mode), 7210 SAS-Sx 10/100GE (standalone mode), and 7210 SAS-T (network mode and access-uplink mode).

See [Queue management policies](#) for more information about configuring slope policies on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

17.1 Overview

Slope policies define WRED parameters used to manage congestion at queuing points in the network.

17.1.1 Configuration guidelines

This section provides configuration guidelines for slope QoS policies.

17.1.2 WRED support on 7210 SAS-T access-uplink and network mode, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE

Two WRED slopes are supported per queue, one each for in-profile or high-priority traffic and out-of-profile or low-priority traffic.

The hardware supports a limited amount of profiles, out of which some are reserved for system internal use and the rest is available for user configuration. It is not possible to allocate a unique profile for each and every queue available on platform. Multiple queues will need to share the same WRED profile. Software manages the allocation of hardware WRED profiles based on user configuration. It automatically allocates a single WRED hardware profile if multiple queues use the same slope parameters (that is, max-average, start-average, drop probability and time average factor). Only if these parameters differ, it allocates a different hardware WRED profile for use by the queue. A WRED profile (that is, each high-slope and low-slope) allows the user to specify the slope parameters, such as maximum average, start average, drop probability, and time average factor (TAF).



Note:

Though the WRED profile is shared among queues, the WRED state (for example, average queue size) is maintained independently for each queue in the hardware.

17.1.3 Basic configurations

A basic slope QoS policy must conform to the following:

- Each slope policy must have a unique policy ID.
- High slope, low slope and non-TCP slope are shut down (default).

- Default values can be modified but parameters cannot be deleted.

17.1.3.1 Create a slope QoS policy

Configuring and applying slope policies is optional. If no slope policy is explicitly applied to a port, a default slope policy is applied.

To create a new slope policy, define the following:

- **a slope policy ID value**

The system will not dynamically assign a value.

- **include a description**

The description provides a brief overview of policy features.

- The high slope for the high priority Random Early Detection (RED) slope graph.
- The low slope for the low priority Random Early Detection (RED) slope graph.
- The time average factor (TAF), a weighting exponent used to determine the portion of the shared buffer instantaneous utilization and shared buffer average utilization used to calculate the new shared buffer average utilization.

Use the following CLI syntax to configure a slope policy for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE:

```
config>qos
  slope-policy name
    description description-string
    high-slope
      start-avg percent
      max-avg percent
      max-prob percent
      no shutdown
    low-slope
      start-avg percent
      max-avg percent
      max-prob percent
      no shutdown
    time-average-factor taf
```

The following displays the slope policy configuration for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE:

Example

```
*A:hw_sass_duth>config>qos>slope-policy# info detail
-----
no description
queue "1"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope
    shutdown
    start-avg 50
    max-avg 75
```



```
        max-prob 75
        exit
        time-average-factor 7
    exit
#
```

17.1.3.2 Applying slope policies

Apply slope policies to the egress buffer pool on the access and network ports.

17.1.3.2.1 Ports

The following CLI syntax examples may be used to apply slope policies to ports:

```
config>port>access>egress>pool>slope-policy name
config>port>network>egress>pool>slope-policy name
```

17.1.4 Default slope policy values

The default access egress and network egress policies are identified as policy-id "default". The default policies cannot be edited or deleted. The following example displays default policy parameters.

Example

```
A:ALA>config>qos# slope-policy default
A:ALA>config>qos>slope-policy# info detail
-----
description "Default slope policy."
queue "1"
    high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
    exit
    low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    time-average-factor 7
exit
queue "2"
    high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
    exit
    low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
```

```
        exit
        time-average-factor 7
    exit
    queue "3"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "4"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "5"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "6"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "7"
```

```

        high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit
queue "8"
    high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit
-----
A:ALA>config>qos>slope-policy#

```

17.2 Service management tasks

17.2.1 Deleting QoS policies

A slope policy is associated by default with 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 ports where it is applied or if the policies are using the slope-policy.

```

ALA-7>config>qos# no slope-policy slopePolicy1
MINOR: QOS #1902 Slope policy has references
ALA-7>config>qos#

```

17.2.1.1 Ports

The following CLI syntax examples can be used to remove slope policies from ports:

```

config>port>access>egress>pool# no slope-policy name
config>port>network>egress>pool# no slope-policy name

```

17.2.1.2 Remove a policy from the QoS configuration

To delete a slope policy, enter the following command:

```
config>qos# no slope-policy policy-id
```

Example:

```
config>qos# no slope-policy slopePolicy1
```

17.2.2 Copying and overwriting QoS policies

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

```
config>qos> copy {slope-policy} source-policy-id dest-policy-id [overwrite]
```

Output example

The following output displays the copied policies.

```
A:ALA-7210M>config>qos#
-----
...
    description "Default slope policy."
    queue "1"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "2"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "3"
```

```
        high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit
queue "4"
    high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit
queue "5"
    high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit
...
-----
A:ALA-7210M>config>qos#
```

17.2.3 Editing QoS policies

You can change existing policies and entries in the CLI or NMS. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors, copy the policy to a work area, make the edits, and then write over the original policy.

17.3 Slope QoS policy command reference

17.3.1 Command hierarchies

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

17.3.1.1 Configuration commands

```
config
- qos
  - [no] slope-policy name
    - description description-string
    - no description
    - queue queue-id
      - [no] high-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
```

17.3.1.2 Operational commands

```
config
- qos
  - copy slope-policy src-name dst-name [overwrite]
```

17.3.1.3 Show commands

```
show
- qos
  - slope-policy [slope-policy-name] [detail]
```

17.3.2 Command descriptions

17.3.2.1 Configuration commands

17.3.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>slope-policy

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

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

description-string

Specifies 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.

17.3.2.1.2 Slope policy QoS commands

slope-policy

Syntax

[no] **slope-policy** *name*

Context

config>qos

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command configures a QoS slope policy.

Default

slope-policy "default"

Parameters

name

Specifies the name of the slope policy. 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.

17.3.2.1.3 Slope policy QoS policy commands

queue

Syntax

queue *queue-id*

Context

config>qos>slope-policy

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

Commands in this context configure the high-priority, low-priority, and non-tcp slope parameters per queue.

Parameters

queue-id

Specifies the ID of the queue for which the drop-rate is to be configured.

Values 1 to 8

high-slope

Syntax

[no] **high-slope**

Context

config>qos>slope-policy>queue

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command defines the high priority Random Early Detection (RED) slope graph. Each buffer pool supports a high priority RED slope for managing access to the shared portion of the buffer pool for high priority or in-profile packets.

The **high-slope** parameters can be changed at any time and the affected buffer pool high priority RED slopes will be adjusted appropriately.

The **no** form of this command reverts the high slope configuration commands to the default values. If the commands within **high-slope** are set to the default parameters, the **high-slope** node will not appear in **save config** and **show config** output unless the **detail** parameter is present.

low-slope

Syntax

[no] **low-slope**

Context

config>qos>slope-policy
config>qos>slope-policy>queue

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command defines the low priority Random Early Detection (RED) slope graph. Each buffer pool supports a low priority RED slope for managing access to the shared portion of the buffer pool for low priority or out-of-profile packets.

The **low-slope** parameters can be changed at any time and the affected buffer pool low priority RED slopes must be adjusted appropriately.

The **no** form of this command reverts 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 *value*

no time-average-factor

Context

config>qos>slope-policy>queue

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command configures 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 Time Average Factor (TAF) value applies to all high, low priority, and non-tcp packets WRED slopes for egress access and network buffer pools controlled by the slope policy.

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

Default

7

Parameters

value

Specifies the 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

17.3.2.1.4 RED slope commands

max-avg

Syntax

max-avg *percent*

no max-avg

Context

config>qos>slope-policy>queue>high-slope

config>qos>slope-policy>queue>low-slope

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command configures the low priority or high priority or non-tcp Weighted Random Early Detection (WRED) slope position for the reserved and 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 reverts the **max-avg** value to the default. If the current **start-avg** setting is larger than the default, an error will occur and the **max-avg** value will not be changed to the default.

Default

max-avg 90 — High slope default is 90% buffer utilization before discard probability is 1

max-avg 75 — Low slope default is 75% buffer utilization before discard probability is 1

Parameters

percent

Specifies the percentage of the reserved and shared buffer space for the buffer pool at which point the drop probability becomes 1. The value entered must be greater or equal to the current setting of **start-avg**. If the entered value is smaller than the current value of **start-avg**, an error will occur and no change will take place.

Values 0 to 100

max-prob

Syntax

max-prob *percent*

no max-prob

Context

config>qos>slope-policy>queue>high-slope

config>qos>slope-policy>queue>low-slope

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command configures the low priority or high priority Random Early Detection (RED) slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one. The *percent* parameter is expressed as a percentage of packet discard probability where always discard is a probability of 1. A **max-prob** value of 80 represents 80% of 1, or a packet discard probability of 0.8.

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

Default

max-prob 80

Parameters

percent

Specifies the maximum drop probability percentage corresponding to the **max-avg**, expressed as a decimal integer.

Values 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 25, 50, 75, 100

shutdown

Syntax

[no] shutdown

Context

config>qos>slope-policy>high-slope

config>qos>slope-policy>low-slope

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

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

start-avg

Syntax

start-avg *percent*

no start-avg

Context

config>qos>slope-policy>queue>high-slope

config>qos>slope-policy>queue>low-slope

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command configures the low priority or high priority Random Early Detection (RED) slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero. The *percent* parameter is expressed as a percentage of the shared buffer size.

The **no** form of this command reverts the **start-avg** value to the default. If the **max-avg** value is smaller than the default, an error will occur and the **start-avg** value will not be changed to the default.

Default

max-avg 70 — High slope default is 70% buffer utilization

max-avg 50 — Low slope default is 50% buffer utilization

Parameters

percent

Specifies the percentage of reserved and shared buffer space for the buffer pool at which the drop starts. The value entered must be lesser 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

17.3.2.1.5 Slope policy QoS policy commands

time-average-factor

Syntax

time-average-factor *value*

no time-average-factor

Context

config>qos>slope-policy>queue

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

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, low priority, and non-tcp packets WRED slopes for egress access and network buffer pools controlled by the slope policy.

17.3.2.2 Operational commands

copy

Syntax

copy slope-policy *src-name dst-name* [**overwrite**]

Context

config>qos

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

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

Specifies the source policy ID that the **copy** command will attempt to copy from and specifies the destination policy ID to which the command will copy a duplicate of the policy. Indicates that the source policy ID and the destination policy ID are slope policy IDs.

overwrite

Keyword to overwrite the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy.

17.3.2.3 Show commands

slope-policy

Syntax

slope-policy [*slope-policy-name*] [**detail**]

Context

show>qos

Platforms

7210 SAS-T (in access-uplink mode and Network), 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE

Description

This command displays slope policy information.

Parameters

slope-policy-name

Specifies the name of the slope policy.

detail

Displays detailed information about the slope policy.

Output

The following output is an example of QoS slope policy information, and [Table 76: Output fields: QoS slope policy](#) describes the output fields.

Sample output

```
*A:SN12345678# show qos slope-policy 100
=====
QoS Slope Policy
```

```

=====
Policy      : 100
Description  : Slope policy 100
-----
Utilization      State      Start-Threshold
-----
High Slope
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1        Down        70           90           75
Queue2        Down        70           90           75
Queue3        Down        70           90           75
Queue4        Down        70           90           75
Queue5        Down        70           90           75
Queue6        Down        70           90           75
Queue7        Down        70           90           75
Queue8        Down        70           90           75
-----
Low Slope
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1        Down        50           75           75
Queue2        Down        50           75           75
Queue3        Down        50           75           75
Queue4        Down        50           75           75
Queue5        Down        50           75           75
Queue6        Down        50           75           75
Queue7        Down        50           75           75
Queue8        Down        50           75           75
-----
Time Avg Factor
-----
Queue Id      Time Avg Factor
-----
Queue1         7
Queue2         7
Queue3         7
Queue4         7
Queue5         7
Queue6         7
Queue7         7
Queue8         7
=====
*A:SN12345678# show qos slope-policy 100 detail

*A:SN12345678#
=====
QoS Slope Policy
=====
Policy      : 100
Description  : Slope policy 100
-----
High Slope
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1        Down        70           90           75
Queue2        Down        70           90           75
Queue3        Down        70           90           75
Queue4        Down        70           90           75
Queue5        Down        70           90           75

```


Queue6	Down	70	90	75
Queue7	Down	70	90	75
Queue8	Down	70	90	75

Low Slope				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	50	75	75
Queue2	Down	50	75	75
Queue3	Down	50	75	75
Queue4	Down	50	75	75
Queue5	Down	50	75	75
Queue6	Down	50	75	75
Queue7	Down	50	75	75
Queue8	Down	50	75	75

Time Avg Factor				

Queue Id	Time Avg Factor			

Queue1	7			
Queue2	7			
Queue3	7			
Queue4	7			
Queue5	7			
Queue6	7			
Queue7	7			
Queue8	7			

Associations				

Object Type	Object Id	Application		Pool

Port	1/1/13	Acc-Egr		default
=====				
*A:SN12345678#				

Table 76: Output fields: QoS slope policy

Label	Description
Policy	The ID that uniquely identifies the policy
Description	A string that identifies the policy's context in the configuration file
Time Avg	The weighting between the previous shared buffer average utilization result and the new shared buffer utilization
Slope Parameters	
Start Avg	Specifies the low priority or high priority RED slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero
Max Avg	Specifies the percentage of the shared buffer space for the buffer pool at which point the drop probability becomes 1, expressed as a decimal integer

Label	Description
Admin State	Up — The administrative status of the RED slope is enabled Down — The administrative status of the RED slope is disabled
Max Prob.	Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one

18 Queue management policies



Note:

This chapter is only applicable to the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

This chapter provides information to configure queue management policies using the command line interface.

18.1 Overview



Note:

The 7210 SAS-R6 and 7210 SAS-R12 equipped with an IMM-c card do not use the configured CBS and MBS parameter values in the queue management policy. The CBS and MBS values are system-defined for each queue. That is, CBS and MBS parameters are not user-configurable. In this section, references to modifying the CBS and MBS values in the queue management policies applies only to the 7210 SAS-R6 and 7210 SAS-R12 equipped with an IMM-b card.

A set of profiles or templates are available in hardware for configuring the queue parameters such as CBS, MBS, and WRED slopes. These profiles are available for use with multiple queues in the system. Queue management policies allow the user to define the queue parameters and allow sharing among the queues.

A single system buffer pool is available for use by all the queues in the system. Users can allocate the amount of buffers that each queue can use by specifying the CBS and MBS parameters in the queue management policy.

Weighted Random Early Detection (WRED) is available to manage buffers during periods of congestion. WRED slopes are supported for each queue in the system. Queue management policies allow the user to configure slope parameters that dictate a WRED profile for each queue. Each queue supports the following slopes:

- slope for in-profile or high priority traffic
- slope for out-of-profile or low priority traffic

Each slope allows specifying the start-average, the max-average, the drop-probability and the Time Average Factor (TAF). Each queue has a default slope policy. Multiple queues in the system can share a single policy. If a policy is shared the system computes the WRED drop probabilities for each of the queues separately based on their average queue length.

18.1.1 Basic configurations

A basic queue management policy must conform to the following restrictions:

- Each slope policy must have a unique policy ID.
- High slope and low slope are shut down by default.
- Default values can be modified but parameters cannot be deleted.

18.1.2 Service management tasks

18.1.2.1 Creating a queue management policy

To create a new queue management policy, define the following:

- a queue management policy name
- a brief description of the policy features
- CBS and MBS values
- high slope for the high priority WRED slope graph
- low slope for the low priority WRED slope graph
- time average factor (TAF)
- slope parameters, such as **max-avg**, **start-avg**, **max-prob**, **time-average-factor**

Use the following CLI syntax to configure a queue management policy.

```
config>qos
  queue-mgmt name
    description description-string
    cbs kbytes
    mbs kbytes
    high-slope
      start-avg percent
      max-avg percent
      max-prob percent
      no shutdown
    low-slope
      start-avg percent
      max-avg percent
      max-prob percent
      no shutdown
    time-average-factor taf
```

Output example

The following is a sample queue management policy configuration output:

```
A:7210>config>qos>queue-mgmt# info
```

```
-----
      high-slope
        shutdown
        start-avg 40
        max-avg 50
      exit
      low-slope
        shutdown
        start-avg 40
        max-avg 80
      exit
      cbs 5000
      mbs 800000
      time-average-factor 7
-----
```

18.1.2.2 Editing QoS policies

About this task

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors, perform the following:

Procedure

Step 1. Copy the policy to a work area.

Step 2. Edit the policy.

Step 3. Overwrite the original policy.

18.2 Queue management policy command reference

18.2.1 Command hierarchies

- [Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#)
- [Operational commands](#)
- [Show commands](#)

18.2.1.1 Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

```
config
- qos
  - queue-mgmt name [create]
  - no queue-mgmt name
    - description description-string
    - no description
    - cbs size-in-kbytes
    - no cbs
    - mbs size-in-kbytes
    - no mbs
    - [no] high-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
- scope {exclusive | template}
- no scope
```

18.2.1.2 Operational commands

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

18.2.1.3 Show commands

```
show
- qos
- queue-mgmt [name] [detail]
```

18.2.2 Command descriptions

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

18.2.2.1 Configuration commands

- [Generic commands](#)
- [Queue management policy QoS commands](#)
- [WRED slope commands](#)

18.2.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>queue-mgmt

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

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

description-string

Specifies 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.

18.2.2.1.2 Queue management policy QoS commands

cbs

Syntax

cbs *size-in-kbytes*

no cbs

Context

config>qos>queue-mgmt

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command specifies the CBS value. The CBS is the minimum depth of the queue.



Note:

The 7210 SAS-R6 and 7210 SAS-R12 equipped with an IMM-c card ignore the CBS parameters configured in the queue management policy. The CBS value is system-defined for each queue and is not user-configurable.

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

Parameters

size-in-kbytes

Specifies the minimum depth of the queue, in kilobytes.

Values 0 to 145000, **default**

high-slope

Syntax

[no] **high-slope**

Context

config>qos>queue-mgmt

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures the in-profile WRED slope parameters.

The **no** form of this command reverts the high-slope configuration commands to the default values.

low-slope

Syntax

[no] **low-slope**

Context

config>qos>queue-mgmt

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures the out-of-profile WRED slope parameters.

The **no** form of this command reverts the low-slope configuration commands to the default values.

mbs

Syntax

mbs *size-in-kbytes*

no mbs

Context

config>qos>queue-mgmt

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command specifies the MBS value. The MBS is the maximum depth of the queue.



Note:

The 7210 SAS-R6 and 7210 SAS-R12 equipped with an IMM-c card ignore the MBS parameters configured in the queue management policy. The MBS value is system-defined for each queue and is not user-configurable.

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

Parameters

size-in-kbytes

Specifies the maximum depth of the queue, in kilobytes.

Values 0 to 145000, **default**

queue-mgmt

Syntax

queue-mgmt *name* [**create**]

no queue-mgmt *name*

Context

config>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures a QoS queue management policy. A set of profiles and templates are available in hardware for configuring the queue parameters such as CBS, MBS, and WRED slopes parameters per queue. These profiles are available for use with multiple queues in the system. Queue management policy allows the user to define the queue parameters and allow for sharing among the queues.

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

Parameters

name

Specifies the name of the queue management policy. 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.

create

Keyword used to create a QoS queue management policy.

time-average-factor

Syntax

time-average-factor *value*

no time-average-factor

Context

config>qos>queue-mgmt

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures the time-average factor (TAF).

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, where a value of zero means the shared buffer instantaneous utilization is used exclusively. A high value weights the new shared buffer average utilization calculation more to the previous shared buffer average utilization value.

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

Default

7

Parameters

value

Specifies the TAF, expressed as a decimal integer.

Values 0 to 15

scope

Syntax

scope {**exclusive** | **template**}

no scope

Context

config>qos>queue-mgmt

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

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

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

Default

template

Parameters

exclusive

Specifies that 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.

template

Specifies that the policy can be applied to multiple interface ports on the router. Default QoS policies are configured with **template** scope. An error is generated if you try to modify the scope parameter from **template** to **exclusive** scope on default policies.

18.2.2.1.3 WRED slope commands

max-avg

Syntax

max-avg *percent*

no max-avg

Context

config>qos>queue-mgmt>high-slope

config>qos>queue-mgmt>low-slope

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures the maximum average value.

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

Default

max-avg 90 — high slope default is 90% buffer utilization

max-avg 75 — low slope default is 75% buffer utilization

Parameters

percent

Specifies the maximum average for the high or low slopes.

Values 0 to 100

max-prob

Syntax

max-prob *percent*

no max-prob

Context

config>qos>queue-mgmt>high-slope

config>qos>queue-mgmt>low-slope

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command configures the maximum probability value.

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

Default

max-avg 75 — high slope default is 75% maximum drop probability corresponding to max-avg

max-avg 75 — low slope default is 75% maximum drop probability corresponding to max-avg

Parameters

percent

Specifies the maximum probability for the high or low slopes.

Values 1 to 99

shutdown

Syntax

[no] shutdown

Context

```
config>qos>queue-mgmt>high-slope  
config>qos>queue-mgmt>low-slope
```

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command enables or disables the administrative status of the WRED slope.

By default, all slopes are shutdown and have to be explicitly enabled (**no shutdown**), which implies a zero drop probability.

The **no** form of this command administratively enables the WRED slope.

Default

shutdown

start-avg

Syntax

```
start-avg percent  
no start-avg
```

Context

```
config>qos>queue-mgmt>high-slope  
config>qos>queue-mgmt>low-slope
```

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command is used to configure the starting average value.

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

Default

max-avg 70 — high slope default is 70% buffer utilization
max-avg 50 — low slope default is 50% buffer utilization

Parameters

percent

Specifies the starting average for the high or low slopes.

Values 0 to 100

18.2.2.2 Operational commands

copy

Syntax

copy queue-mgmt *src-name* *dst-name* [**overwrite**]

Context

config>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command copies the existing queue management policy entries to another queue management 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.

Parameters

src-name

Specifies the name of the source policy, up to 32 characters.

dst-name

Specifies the name of the destination policy, up to 32 characters.

overwrite

Keyword to overwrite the information in the destination policy by the information in the source policy.

18.2.2.3 Show commands

queue-mgmt

Syntax

queue-mgmt [*name*] [**detail**]

Context

show>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command displays queue management policy information.

Parameters

- name**
Specifies the name of the queue management policy.
- detail**
Displays detailed information about the queue management policy.

Output

The following output is an example of queue management policy information, and [Table 77: Output fields: queue management policy](#) describes the output fields.

Sample output for 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, and 7210 SAS-Mxp

```
*A:Dut-A>config# /show qos queue-mgmt 100 detail
=====
QoS Queue Management Policy
=====
Policy          : 100
Description      : (Not Specified)
CBS             : 100                      MBS             : 200
Time Avg        : 7
High Slope Parameters
-----
Start Avg       : 70                      Admin State      : Enabled
Max Avg         : 90                      Max Prob.        : 75
Low Slope Parameters
-----
Start Avg       : 50                      Admin State      : Enabled
Max Avg         : 75                      Max Prob.        : 75
Associations
SAP Egress
-----
No SAP Egress Associations found.
-----
Network Queues
-----
No Network Queue Policy Associations found.
-----
Access Egress
-----
Access Egress Policy Id      : 199
Queue Ids                   : 1, 2, 3, 4, 5, 6, 7, 8
-----
Access-Port Associations
-----
Port-id : 1/1/6
-----
Access Egress Queue Override
-----
No Access Egress queue override Associations found.
-----
=====
*A:Dut-A>config# /show qos queue-mgmt 200 detail
=====
QoS Queue Management Policy
=====
```

```

Policy      : 200
Description : (Not Specified)
CBS         : 125
Time Avg    : 7
MBS         : 500
High Slope Parameters
-----
Start Avg   : 70
Max Avg     : 90
Admin State : Enabled
Max Prob.   : 75
Low Slope Parameters
-----
Start Avg   : 50
Max Avg     : 75
Admin State : Disabled
Max Prob.   : 75
Associations
SAP Egress
-----
No SAP Egress Associations found.
-----
Network Queues
-----
Network Queue Policy Name : 200
Queue Ids                 : 1, 2, 3, 4, 5, 6, 7, 8
-----
Network-Port Associations
-----
Port-id : 1/1/26
-----
Access Egress
-----
No Access Egress Associations found.
-----
Access Egress Queue Override
-----
Port-id : 1/1/6   Queue Ids : 1
-----
=====

```

Table 77: Output fields: queue management policy

Label	Description
Policy	Displays the ID that uniquely identifies the policy
Description	Displays a string that identifies the policy's context in the configuration file
Time Avg	Displays the weighting between the previous shared buffer average utilization result and the new shared buffer utilization
CBS	Displays the committed burst size
MBS	Displays the maximum burst size
Slope Parameters	
Start Avg	Displays, in percentage, the low priority or high priority RED slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero

Label	Description
Max Avg	Displays the percentage of the shared buffer space for the buffer pool at which point the drop probability becomes 1, expressed as a decimal integer
Admin State	Up — The administrative status of the RED slope is enabled Down — The administrative status of the RED slope is disabled
Max Prob.	Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one
Service Associations	
SAP Egress Policy Id	Displays the SAP egress policy ID
Queue Ids	Displays the queue IDs
Access Egress Associations	
Access Egress Policy Id	Displays the access-egress policy ID
Queue Ids	Displays the queue IDs
Access-Port Associations	
Port-id	Displays the port ID

19 Remark policies

This section provides information to configure remark policies using the command line interface. This section is applicable to only 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T devices in network mode.


19.1 Overview

The remark policies are used to configure the marking behavior for the system at the egress of access SAP access port, network ports, hybrid ports, and network IP interfaces on network ports or hybrid ports (for support available on different 7210 SAS platforms, see [Table 78: Summary of remark policy and attachment points for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE \(network mode\)](#) and [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#)). These policies allow the user to define the forwarding class to egress marking values and allow them to use the available hardware resources efficiently. Based on the packet encapsulation used, the remark policy allows the user to define and associate appropriate policies to service egress, and network QoS policies. The 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T support the use of the following types of remark policies for different QoS policies:

- **dot1p**
This policy is used for service egress, access port egress, and network QoS (**port** type policies).
- **dscp**
This policy is used for access port egress marking and network QoS policies of the **port** type.
- **lsp-exp**
This policy is used for network QoS policies of the **ip-interface** type.
- **dot1p-dscp**
This policy is used for access port egress marking and network QoS policies of the **port** type.
- **dot1p-lsp-exp-shared**
This policy is used for access port egress, access port egress marking, and network QoS policies of the **ip-interface** type.

The type of the remark policy identifies the bits marked in the packet header. Each of these remark policy types can be associated with only appropriate QoS policies and service entities as listed in the following tables.

Table 78: Summary of remark policy and attachment points for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE (network mode)

Remark policy type	QoS policy	Attachment point	Packet header bits marked
dot1p	access-egress policy	Access port	<ul style="list-style-type: none"> Dot1p bits in L2 Ethernet header for service packets sent out of all access SAPs configured on access ports.
	network policy (type port)	Network Port and Hybrid Port	<ul style="list-style-type: none"> Dot1p bits in the L2 Ethernet header for IP packets sent out of network port and hybrid port.
dscp	access-egress policy	Access port	<ul style="list-style-type: none"> IP DSCP bits in the IP header for service packets sent out of access SAPs configured on access ports.
	network policy (type port)	Network Port and Hybrid Port	<ul style="list-style-type: none"> IP DSCP bits in the IP header for IP packets sent out of network port and hybrid port. <div>  Note: For the IP traffic sent out of L2 SAPs on hybrid ports, the DSCP bits are not marked. </div>
lsp-exp	network policy (type ip-interface)	Network IP interface	<ul style="list-style-type: none"> EXP bits in the MPLS header for MPLS packets sent out of IP interface configured on network port and hybrid port.
dot1p-lsp-exp-shared	access-egress policy	Access port	<ul style="list-style-type: none"> Dot1p bits in L2 Ethernet header for service packets sent out of access SAPs configured on access ports.
	network policy (type IP-interface)	Network IP interface	<ul style="list-style-type: none"> EXP bits in the MPLS header for MPLS packets sent out of IP interface configured on network port and hybrid port. Dot1p bits in the L2 Ethernet header for MPLS packets sent out of network and hybrid port.
dot1p-dscp	access-egress policy	Access port	<ul style="list-style-type: none"> Dot1p bits in L2 Ethernet header and IP DSCP bits in the IP


Remark policy type	QoS policy	Attachment point	Packet header bits marked
			header for service packets sent out of access SAPs configured on access ports.
	network policy (type port)	Network Port and Hybrid Port	<ul style="list-style-type: none"> Dot1p bits in the L2 Ethernet header for IP packets and IP DSCP bits in the IP header for IP packets sent out of network port and hybrid port.  <p>Note: For the IP traffic sent out of L2 SAPs on hybrid ports, the DSCP bits are not marked.</p>

Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Remark policy type	Qos policy	Attachment point	Packet header bits marked
dot1p	SAP-egress policy	Service/SAP egress	Dot1p bits in the L2 Ethernet header for service packets sent out of a SAP egress
	Access-egress policy	Access port	Dot1p bits in L2 Ethernet header for service packets sent out of all access SAPs configured on access ports ¹⁸
	Network policy (port type)	Network port and hybrid port	Dot1p bits in the L2 Ethernet header for IP packets sent out of network port and hybrid port
dscp	Access egress policy	Access port	IP DSCP bits in the IP header for service packets sent out of access SAPs configured on access ports ¹⁸
	Network policy (port type)	Network port and hybrid port	IP DSCP bits in the IP header for IP packets sent out of network port and hybrid port ¹⁸
lsp-exp	Network policy (IP-interface type)	Network IP interface	EXP bits in the MPLS header for MPLS packets sent out of IP interface configured on network port and hybrid port

¹⁸ For the IP traffic sent out of L2 SAPs on hybrid ports, the DSCP bits are not marked.

Remark policy type	Qos policy	Attachment point	Packet header bits marked
dot1p-lsp-exp-shared	Access egress policy	Access port	Dot1p bits in L2 Ethernet header for service packets sent out of access SAPs configured on access ports ¹⁸
	Network policy (IP-interface type or port type)	Network IP interface or network port	EXP bits in the MPLS header for MPLS packets sent out of IP interface configured on network port and hybrid port Dot1p bits in the L2 Ethernet header for MPLS packets sent out of network port and hybrid port
	SAP egress policy	Service/SAP egress	Dot1p bits in the L2 Ethernet header for service packets sent out of a SAP egress
dot1p-dscp	Access egress policy	Access port	Dot1p bits in L2 Ethernet header and IP DSCP bits in the IP header for service packets sent out of access SAPs configured on access ports ¹⁸
	Network policy (port type)	Network port and hybrid port	Dot1p bits in the L2 Ethernet header for MPLS and IP packets and IP DSCP bits in the IP header for IP packets sent out of network port and hybrid port ¹⁸

19.1.1 Configuration guidelines

- For access port and access SAP marking functionality (SAP-based and port-based), the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 support SAP-based egress marking and port-based egress marking on only access ports when using access SAP-based egress queuing. Users have an option to use either SAP-based marking or port-based marking. In SAP-based marking, the remark policy defined in the SAP egress policy associated with each SAP is used to mark the packets egressing out of the SAP, if marking is enabled. In port-based marking, the remark policy defined in the access-egress policy associated with the access port determines the marking values to use for all the SAPs defined on that port. See [Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information.
- The 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 support only access port-based marking when port-based egress queues are configured. See [Access egress QoS policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information.
- If DSCP remarking or both dot1p and DSCP remarking is specified, the DSCP field for traffic sent out of SAPs configured in a Layer 2 services are not marked.

19.1.2 Basic configurations

A basic remark policy must confirm to the following.

- Each remark policy must have a unique policy ID.
- The remark policy type must be specified.
- The forwarding class to egress marking values must be specified.

19.1.2.1 Creating a remark policy

To create a new remark policy, define the following:

- A remark policy name and type is optional and by default it is 'dot1p'.
- Provide a brief description of the policy features.
- Specify the forwarding class to egress marking values.

Use the following syntax to configure a remark policy.

```
A:7210-T>config>qos# remark 122 remark-type dscp create
```

Output example

The following is a sample remark policy configuration output.

```
*A:7210SAS>config>qos>remark# info detail
-----
description "Default Remarking Policy for dot1P and DSCP"
fc af
  dscp-in-profile af11
  dscp-out-profile af12
  no lsp-exp-in-profile
  no lsp-exp-out-profile
  no dot1p-lsp-exp-in-profile
  no dot1p-lsp-exp-out-profile
  no de-mark
  no dot1p
  dot1p-in-profile 3
  dot1p-out-profile 2
exit
fc be
  dscp-in-profile be
  dscp-out-profile be
  no lsp-exp-in-profile
  no lsp-exp-out-profile
  no dot1p-lsp-exp-in-profile
  no dot1p-lsp-exp-out-profile
  no de-mark
  no dot1p
  dot1p-in-profile 0
  dot1p-out-profile 0
exit
fc ef
  dscp-in-profile ef
  dscp-out-profile ef
  no lsp-exp-in-profile
  no lsp-exp-out-profile
  no dot1p-lsp-exp-in-profile
```

```
        no dotlp-lsp-exp-out-profile
        no de-mark
        no dotlp
        dotlp-in-profile 5
        dotlp-out-profile 5
    exit
    fc h1
        dscp-in-profile nc1
        dscp-out-profile nc1
        no lsp-exp-in-profile
        no lsp-exp-out-profile
        no dotlp-lsp-exp-in-profile
        no dotlp-lsp-exp-out-profile
        no de-mark
        no dotlp
        dotlp-in-profile 6
        dotlp-out-profile 6
    exit
    fc h2
        dscp-in-profile af41
        dscp-out-profile af41
        no lsp-exp-in-profile
        no lsp-exp-out-profile
        no dotlp-lsp-exp-in-profile
        no dotlp-lsp-exp-out-profile
        no de-mark
        no dotlp
        dotlp-in-profile 4
        dotlp-out-profile 4
    exit
    fc l1
        dscp-in-profile af21
        dscp-out-profile af22
        no lsp-exp-in-profile
        no lsp-exp-out-profile
        no dotlp-lsp-exp-in-profile
        no dotlp-lsp-exp-out-profile
        no de-mark
        no dotlp
        dotlp-in-profile 3
        dotlp-out-profile 2
    exit
    fc l2
        dscp-in-profile cs1
        dscp-out-profile cs1
        no lsp-exp-in-profile
        no lsp-exp-out-profile
        no dotlp-lsp-exp-in-profile
        no dotlp-lsp-exp-out-profile
        no de-mark
        no dotlp
        dotlp-in-profile 1
        dotlp-out-profile 1
    exit
    fc nc
        dscp-in-profile nc2
        dscp-out-profile nc2
        no lsp-exp-in-profile
        no lsp-exp-out-profile
        no dotlp-lsp-exp-in-profile
        no dotlp-lsp-exp-out-profile
        no de-mark
        no dotlp
        dotlp-in-profile 7
```

```
dot1p-out-profile 7
exit
-----
*A:7210SAS>config>qos>remark#
```

19.1.2.2 Editing QoS policies

About this task

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors perform the following:

Procedure

- Step 1.** Copy the policy to a work area.
- Step 2.** Edit the policy.
- Step 3.** Overwrite the original policies.

19.2 Remark policy command reference

19.2.1 Command hierarchies

- [Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T in network mode](#)
- [Operational commands](#)
- [Show commands](#)

19.2.1.1 Configuration commands for 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T in network mode

```
config
- qos
  - [no] remark policy-id [remark-type remarking-type] [create]
  - [no] description description-string
  - [no] fc fc-name
    - [no] de-mark [force de-value]
    - 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
- dot1p-lsp-exp-in-profile dot1p | lsp-exp value
- no dot1p-lsp-exp-in-profile
- dot1p-lsp-exp-out-profile dot1p | lsp-exp value
- no dot1p-lsp-exp-out-profile
```

19.2.1.2 Operational commands

```
config
- qos
- copy remark src-pol dst-pol [overwrite]
```

19.2.1.3 Show commands

```
show
- qos
- remark-policy [policy-id] [association | detail]
```

19.2.2 Command descriptions

19.2.2.1 Configuration commands

19.2.2.1.1 Generic commands

description

Syntax

[no] **description** *description-string*

Context

config>qos>remark

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

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

description-string

Specifies 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.

19.2.2.1.2 Remark policy QoS commands

remark

Syntax

remark *policy-id* [**remark-type** **marking-type**] [**create**]

no remark *policy-id*

Context

config>qos

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command creates a new remark policy of the specified type.

The following types of remark policies are available:

- dot1p
- dscp
- dot1p-dscp
- lsp-exp
- dot1p-lsp-exp-shared

The **remark-type** of the policy also determines the values user is allowed to configure in the policy and also the QoS policy with which this remark policy can be associated with. See [Table 78: Summary of remark policy and attachment points for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE \(network mode\)](#) and [Table 79: Summary of remark policy and attachment points for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12](#) for more information about remark policy.

Parameters

policy-id

Specifies the policy ID of the remark policy.

Values 1 to 65535

remarking-type

Specifies the type of marking values in the remark policy.

Values dot1p — Specify the FC to 802.1 dot1p value to use for marking. It is the default used if user does not explicitly specify the **remarking-type** value.

 dscp — Specify the FC to IP DSCP value to use for marking.

 dot1p-dscp — Specify FC to both dot1p and IP DSCP values to use for marking. lsp-exp — Specify the FC to MPLS EXP values to use for marking.

 dot1p-lsp-exp-shared — Specify FC to MPLS EXP and dot1p values to use for marking.

 These policies share a common resource in hardware and a single FC is mapped to the same MPLS EXP value and dot1p value.

fc

Syntax

[no] fc *fc-name*

Context

config>qos>remark

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command specifies the forwarding class name and provides the context to configure the marking value for the FC. Based on the type of remark policy created, the **fc** command allows the user to specify the appropriate marking values. The **fc** command overrides the default parameters for the forwarding class to the values defined.

The **no** form of this command removes the forwarding class to marking values map associated with the FC. The FC reverts to the defined parameters in the default remark policy.

Parameters

fc-name

Specifies a case-sensitive system-defined forwarding class name for which policy entries are created.

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

19.2.2.1.3 Remark policy forwarding class commands

de-mark

Syntax

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

Context

config>qos>remark-policy>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

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

If no *de-value* is present, the default values are used for the marking of the DE bit, as defined in the IEEE 802.1ad-2005 standard. For example 0 for in-profile packets, 1 for out-of-profile ones.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS, the command dictates the marking of the DEI bit for both the BVID and ITAG.

If this command is not used, the DEI bit should be preserved if an ingress TAG exist or set to zero otherwise.

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

Parameters

de-value

Specifies the DEI value to use for this forwarding class.

Values 0 or 1

dot1p

Syntax

[no] dot1p *dot1p-value*

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the **dot1p** command has no effect.

DEI marking must be enabled using the **de-mark** command along with this command for the command to take effect. When **de-mark** command is configured along with this command, then the DEI bit is marked in the packet to indicate the profile of the packet. The DEI bit is marked to 0 to indicate in-profile/green packet and 1 to indicate out-of-profile/yellow packet. If the **force de-value** parameter is specified then the DEI bit is set to specified value for all packets.

If the **no** form of this command is executed then software uses the dot1p-in-profile and dot1p-out-profile if configured; otherwise, the software uses the default values.

Default

no dot1p

Parameters

dot1p-value

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

dot1p-in-profile

Syntax

dot1p-in-profile *dot1p-priority*

no dot1p-in-profile

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command allows to mark on an egress the in and out of profile status through a certain dot1p combination, similarly with the DEI options. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with in-profile status (or green color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the **dot1p** command has no effect.

If DEI marking is enabled using the **de-mark** command and the command **dot1p dot1p-value** is used to configure the dot1p value, then this command has no effect. In other words, enabling DEI marking has precedence over this command and the system ignores this command.

When this command is used, the DEI Bit is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DEI bit is set to 0.

When the **dot1p dot1p-value** command is used and **de-mark** is enabled, the *dot1p-value* is used for the entire forwarding class. This command is mutually exclusive to the use of the **dot1p** command.

The **no** form of this command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.

Default

0

Parameters

dot1p-priority

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

dot1p-out-profile

Syntax

dot1p-out-profile *dot1p-priority*

no dot1p-out-profile

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command allows to mark on an egress the in and out of profile status via a certain dot1p combination, similarly with the DEI options. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with out-of-profile status (or yellow color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the **dot1p** command has no effect.

If DEI marking is enabled using the **de-mark** command and the *dot1p-value* is configured, then this command has no effect. In other words, enabling DEI marking has precedence over this command and the system ignores this command.

When this command is used the DEI Bit is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DEI bit is set to 0.

When the command **dot1p** *dot1p-value* is used and **de-mark** is enabled, it means that the *dot1p-value* is used for the entire forwarding class. These two variants of the command are mutually exclusive. In other words, this command is mutually exclusive to use of the **dot1p** *dot1p-value* command .

The **no** form of this command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.

Default

0

Parameters

dot1p-priority

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

dot1p-inner

Syntax

[no] **dot1p-inner** [in-profile *dot1p-value*] [out-profile *dot1p-value*]

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command allows to mark on an egress the in and out profile status through a certain dot1p combination. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with in-profile status (or green color) of *fc-name* that egress out of QinQ SAP only (for example, SAPs configured with two VLAN tags explicitly defined, such as SAP 1/1/5:10.100) use the explicitly defined in-profile and out-profile *dot1p-value*.

This command has no effect for egress packets sent out of all other non-QinQ SAPs, such as dot1q SAP and null SAP. Additionally, if the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, this command has no effect. In other words, this command takes effect, only when the node adds 2 tags to the packet on the egress.

This variant of the command is mutually exclusive to the use of **dot1p-inner** command. In other words, user has a choice to use either this command or the **dot1p-inner** command but not both together.

If the **no** form of this command is executed, default remarking values are used for marking the inner VLAN.

Default

no dot1p-inner

Parameters

in-profile dot1p-value

Specifies the dot1p value to use for in-profile packets.

Values 0 to 7

out-profile dot1p-value

Specifies the dot1p bits to use for the out-profile packets.

Values 0 to 7

dscp-in-profile

Syntax

dscp-in-profile *dscp-name*

no dscp-in-profile

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command specifies the in-profile DSCP name for the forwarding class. When marking is set, the corresponding DSCP value is used to mark all IP packets with in-profile status, on the egress of this forwarding class queue.

When multiple DSCP names are associated with the forwarding class in the policy, the last name entered will overwrite the previous value.

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

Parameters

dscp-name

Specifies the system- or user-defined, case-sensitive DSCP 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, cp35, 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

dscp-out-profile

Syntax

dscp-out-profile *dscp-name*

no dscp-out-profile

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command specifies the out-of-profile DSCP name for the forwarding class. When marking is set, the corresponding DSCP value is used to mark all IP packets with out-of-profile status, on the egress of this forwarding class queue.

When multiple DSCP names are associated with the forwarding class in the policy, the last name entered will overwrite the previous value.

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

Parameters

dscp-name

Specifies the system- or user-defined, case-sensitive *dscp-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, cp35, 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
---------------	---

lsp-exp-in-profile

Syntax

lsp-exp-in-profile *lsp-exp-value*

no lsp-exp-in-profile

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

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

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

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

Default

- Policy ID 2 — Factory setting
- Policy ID 3 to 65535 — Policy-id setting

Parameters

lsp-exp-value
Specifies 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>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

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

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

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

Default

- Policy ID2 — Factory setting
- Policy ID 3 to 65535 — Policy-id setting

Parameters

mpls-exp-value

Specifies the 3-bit MPLS EXP bit value, expressed as a decimal integer.

Values 0 to 7

lsp-exp-in-profile

Syntax

lsp-exp-in-profile *lsp-exp-value*

no lsp-exp-in-profile

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command specifies the in-profile LSP EXP value for the forwarding class. This value is used for all in-profile LSP labeled packets which require marking the egress on the forwarding class queue.

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

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

Parameters

lsp-exp-value

Specifies a 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>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command specifies the in-profile LSP EXP value for the forwarding class. This value is used for all out-of-profile LSP labeled packets which require marking the egress on the forwarding class queue.

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

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

Parameters

lsp-exp-value

Specifies a 3-bit LSP EXP bit value, expressed as a decimal integer.

Values 0 to 7

dot1p-lsp-exp-in-profile

Syntax

dot1p-lsp-exp-in-profile *dot1p* | *lsp-exp value*

no dot1p-lsp-exp-in-profile

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command specifies the in-profile dot1p LSP EXP value for the forwarding class. This value is used for all in-profile LSP labeled packets which require marking the egress on the forwarding class queue.

When multiple dot1p LSP EXP values are associated with the forwarding class at network egress, the last name entered overwrites the previous value.

The **no** form of this command reverts to the factory default in-profile dot1p LSP EXP value.

Parameters

dot1p | *lsp-exp value*

Specifies a 3-bit dot1p LSP EXP bit value, expressed as a decimal integer.

Values 0 to 7

dot1p-lsp-exp-out-profile

Syntax

dot1p-lsp-exp-out-profile *dot1p* | *lsp-exp value*

no dot1p-lsp-exp-out-profile

Context

config>qos>remark>fc

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command specifies the out-profile dot1p LSP EXP value for the forwarding class. This value is used for all out-of-profile LSP labeled packets which require marking the egress on the forwarding class queue.

When multiple dot1p LSP EXP values are associated with the forwarding class at network egress, the last name entered overwrites the previous value.

The **no** form of this command reverts to the factory default in-profile dot1p LSP EXP value.

Parameters

dot1p | lsp-exp value

Specifies a 3-bit dot1p LSP EXP bit value, expressed as a decimal integer.

Values 0 to 7

19.2.2.2 Operational commands

copy

Syntax

copy remark src-pol dst-pol [overwrite]

Context

config>qos

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command copies existing remark policy entries to another remark 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 the destination policy already exists, the **overwrite** keyword must be specified.

Parameters

- src-pol*

Specifies the source policy.

Values 1 to 65535
- dst-pol*

Specifies the destination policy.

Values 1 to 65535
- overwrite*

Keyword to overwrite the information in the destination policy by the information in the source policy.

19.2.2.3 Show commands

remark-policy

Syntax

remark-policy [*policy-id*] [**association** | **detail**]

Context

show>qos

Platforms

Supported on all 7210 SAS as described in this document, except those operating in access-uplink mode

Description

This command displays remark policy information.

Parameters

- policy-id*

Specifies the ID of the remark policy.
- detail**

Displays detailed information about the remark policy.

Output

The following output is an example of QoS remark policy information, and [Table 80: Output fields: remark policy](#) describes the output fields.

Sample output

```
*A:SAS-C>config>qos# show qos remark-policy
```

```

=====
SAS Remarking Policies
=====
Policy-Id Type Description
-----
1 dscp Default Remarking Policy for DSCP
2 dot1p-lsp-exp-shared Default Remarking Policy for dot1P and LSP*
500 dot1p-lsp-exp-shared
505 dot1p-lsp-exp-shared
510 dot1p-lsp-exp-shared
515 dot1p-lsp-exp-shared
520 dot1p-lsp-exp-shared
525 dot1p-lsp-exp-shared
530 dot1p-lsp-exp-shared
535 dot1p-lsp-exp-shared
540 dot1p-lsp-exp-shared
545 dot1p-lsp-exp-shared
550 dot1p
555 dot1p
560 dot1p
565 dot1p
570 dot1p
=====
•indicates that the corresponding row element may have been truncated.

*A:SAS-C>config>qos# show qos remark-policy 500 association

=====
QoS Remarking Policies
=====
-----
Remark Policy-id : 500 Type : dot1p-lsp-exp-shared
Description : (Not Specified)
-----
Associations
-----
SAP Egress
-----
SAP Egress Policy Id : 5001
-----
Associations
-----
Service-Id : 500 (VPLS) Customer-Id : 1
- SAP : lag-2:500
SAP Egress Policy Id : 5701
-----
Associations
-----
Service-Id : 570 (VPLS) Customer-Id : 1
- SAP : lag-2:570
SAP Egress Policy Id : 6401
-----
Associations
-----
Service-Id : 640 (VPLS) Customer-Id : 1
- SAP : lag-2:640
SAP Egress Policy Id : 10001

```

```
-----
Associations
-----
Service-Id      : 1000 (VPLS)                Customer-Id   : 1
- SAP : lag-2:1000

SAP Egress Policy Id      : 10701

-----
Associations
-----
Service-Id      : 1070 (VPLS)                Customer-Id   : 1
- SAP : lag-2:1070

SAP Egress Policy Id      : 11401

-----
Associations
-----
Service-Id      : 1140 (VPLS)                Customer-Id   : 1
- SAP : lag-2:1140

SAP Egress Policy Id      : 15001

-----
Associations
-----
Service-Id      : 1500 (VPLS)                Customer-Id   : 1
- SAP : lag-4:1500

SAP Egress Policy Id      : 15701

-----
Associations
-----
Service-Id      : 1570 (VPLS)                Customer-Id   : 1
- SAP : lag-4:1570

SAP Egress Policy Id      : 16401

-----
Associations
-----
Service-Id      : 1640 (VPLS)                Customer-Id   : 1
- SAP : lag-4:1640

SAP Egress Policy Id      : 20001

-----

Network
-----
Network Policy Id      : 50

-----
Interface Association
-----
Interface      : ip-192.168.105.4
IP Addr.       : 192.168.105.4/24            Port Id      : 1/1/23
Interface      : ip-192.168.20.4
IP Addr.       : 192.168.20.4/24            Port Id      : lag-3
Interface      : ip-192.168.45.4
```



```
IP Addr.      : 192.168.45.4/24          Port Id      : lag-5
Interface     : ip-192.168.80.4
IP Addr.      : 192.168.80.4/24          Port Id      : 1/1/22
Network Policy Id      : 550

-----
Interface Association
-----
Interface     : ip-192.168.100.4
IP Addr.      : 192.168.100.4/24          Port Id      : 1/1/23
Interface     : ip-192.168.40.4
IP Addr.      : 192.168.40.4/24          Port Id      : lag-5
Interface     : ip-198.162.65.4
IP Addr.      : 192.168.65.4/24          Port Id      : 1/1/13
Network Policy Id      : 1050

-----
Interface Association
-----
No Interface Association Found.

Network Policy Id      : 1550

-----
Interface Association
-----
No Interface Association Found.

Network Policy Id      : 2050

-----
Interface Association
-----
No Interface Association Found.

=====
*A:SAS-C>config>qos# show qos remark-policy 500 detail
=====
QoS Remarking Policies
=====
Remark Policy-id      : 500          Type      : dot1p-lsp-exp-shared
Description           : (Not Specified)

-----
FC Name      dot1P / LSP EXP      dot1P / LSP EXP
              In Value             Out Value
-----
be           3                    6
l2           1                    4
af           6                    1
l1           7                    2
h2           0                    3
ef           5                    0
h1           4                    7
nc           2                    5

-----
Associations
-----
SAP Egress
```

```
-----
SAP Egress Policy Id      : 5001
-----
Associations
-----
Service-Id      : 500 (VPLS)      Customer-Id : 1
- SAP : lag-2:500
-----
SAP Egress Policy Id      : 5701
-----
Associations
-----
Service-Id      : 570 (VPLS)      Customer-Id : 1
- SAP : lag-2:570
-----
SAP Egress Policy Id      : 6401
-----
Associations
-----
Service-Id      : 640 (VPLS)      Customer-Id : 1
- SAP : lag-2:640
-----
SAP Egress Policy Id      : 10001
-----
Associations
-----
Service-Id      : 1000 (VPLS)     Customer-Id : 1
- SAP : lag-2:1000
-----
SAP Egress Policy Id      : 10701
-----
Associations
-----
Service-Id      : 1070 (VPLS)     Customer-Id : 1
- SAP : lag-2:1070
-----
SAP Egress Policy Id      : 11401
-----
Associations
-----
Service-Id      : 1140 (VPLS)     Customer-Id : 1
- SAP : lag-2:1140
-----
SAP Egress Policy Id      : 15001
-----
Associations
-----
Service-Id      : 1500 (VPLS)     Customer-Id : 1
- SAP : lag-4:1500
-----
SAP Egress Policy Id      : 15701
-----
Associations
-----
Service-Id      : 1570 (VPLS)     Customer-Id : 1
- SAP : lag-4:1570
-----
```

```
SAP Egress Policy Id          : 16401
-----
Associations
-----
Service-Id      : 1640 (VPLS)          Customer-Id : 1
- SAP : lag-4:1640

SAP Egress Policy Id          : 20001
-----
Associations
-----
Service-Id      : 2000 (Epipe)         Customer-Id : 1
- SAP : 1/1/2:2000

SAP Egress Policy Id          : 20701
-----
Network
-----
Network Policy Id             : 50
-----
Interface Association
-----
Interface      : ip-192.168.105.4
IP Addr.       : 192.168.105.4/24      Port Id       : 1/1/23
Interface      : ip-192.168.20.4
IP Addr.       : 192.168.20.4/24      Port Id       : lag-3
Interface      : ip-192.168.45.4
IP Addr.       : 192.168.45.4/24      Port Id       : lag-5
Interface      : ip-192.168.80.4
IP Addr.       : 192.168.80.4/24      Port Id       : 1/1/22
Network Policy Id             : 550
-----
Interface Association
-----
Interface      : ip-192.168.100.4
IP Addr.       : 192.168.100.4/24     Port Id       : 1/1/23
Interface      : ip-192.168.40.4
IP Addr.       : 192.168.40.4/24     Port Id       : lag-5
Interface      : ip-192.168.65.4
IP Addr.       : 192.168.65.4/24     Port Id       : 1/1/13
Network Policy Id             : 1050
-----
Interface Association
-----
No Interface Association Found.

Network Policy Id             : 1550
-----
Interface Association
-----
No Interface Association Found.

Network Policy Id             : 2050
-----
Interface Association
```

```
-----
No Interface Association Found.
-----
=====

*A:SAS-C>config>qos# show qos remark-policy 500 detail

=====
QoS Remarking Policies
=====
-----
Remark Policy-id      : 500                Type      : dot1p-lsp-exp-shared
Description           : (Not Specified)
-----

FC Name      dot1P / LSP EXP      dot1P / LSP EXP
              In Value             Out Value
-----
be           3                    6
l2           1                    4
af           6                    1
l1           7                    2
h2           0                    3
ef           5                    0
h1           4                    7
nc           2                    5
-----

Associations
-----
SAP Egress
-----
SAP Egress Policy Id      : 5001
-----

Associations
-----
Service-Id      : 500 (VPLS)                Customer-Id : 1
- SAP : lag-2:500
-----
SAP Egress Policy Id      : 5701
-----

Associations
-----
Service-Id      : 570 (VPLS)                Customer-Id : 1
- SAP : lag-2:570
-----
SAP Egress Policy Id      : 6401
-----

Associations
-----
Service-Id      : 640 (VPLS)                Customer-Id : 1
- SAP : lag-2:640
-----
SAP Egress Policy Id      : 10001
-----

-----
Network
-----
```

```

Network Policy Id          : 50
-----
Interface Association
-----
Interface      : ip-192.168.105.4
IP Addr.       : 192.168.105.4/24      Port Id        : 1/1/23
Interface      : ip-192.168.20.4
IP Addr.       : 192.168.20.4/24      Port Id        : lag-3
Interface      : ip-192.168.45.4
IP Addr.       : 192.168.45.4/24      Port Id        : lag-5
Interface      : ip-192.162.80.4
IP Addr.       : 192.162.80.4/24      Port Id        : 1/1/22
Network Policy Id          : 550
-----
Interface Association
-----
Interface      : ip-192.168.100.4
IP Addr.       : 192.168.100.4/24      Port Id        : 1/1/23
Interface      : ip-192.168.40.4
IP Addr.       : 192.168.40.4/24      Port Id        : lag-5
Interface      : ip-192.168.65.4
IP Addr.       : 192.168.65.4/24      Port Id        : 1/1/13
Network Policy Id          : 1050
-----
Interface Association
-----
No Interface Association Found.

Network Policy Id          : 1550
-----
Interface Association
-----
No Interface Association Found.

Network Policy Id          : 2050
-----
Interface Association
-----
No Interface Association Found.
=====

```

Table 80: Output fields: remark policy

Label	Description
Policy ID	The ID that uniquely identifies the policy
Remark Policy-id	Displays the policy ID of the remark policy
Type	Displays the type of remark policy
Description	A string that identifies the policy's context in the configuration file
FC Name	Specifies the forwarding class name

Label	Description
dot1P/LSP EXP In value	dot1p/LSP EXP value for in-profile packets
dot1P/LSP EXP Out value	dot1p/LSP EXP value for out-of-profile packets
DCSP In value	DSCP value used for in-profile packets
DCSP Out value	DSCP value used for out-of-profile packets
Service Associations	
SAP Egress Policy Id	Displays the policy ID of the SAP Egress policy
Service-Id	The unique service ID number which identifies the service in the service domain
Customer-Id	Specifies the customer ID which identifies the customer to the service
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP ingress policy is applied
Network	
Network Policy Id	Displays the network policy ID
Interface Association	
Interface	Displays the associated interface
IP Addr.	Displays the IP address of the interface

20 Access ingress QoS policies



Note:

This feature is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone).

This chapter provides information about configuring access ingress QoS policies using the command line interface.

20.1 Overview

An access ingress QoS policy is applied to the physical port instead of the SAP. It applies to all SAPs configured on the specific access port. To configure a port-based access ingress QoS policy, the **access-ingress-qos-mode** command must be configured with the **port-mode** option specified.



Note:

See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for more information about the **access-ingress-qos-mode** command.

At ingress, the access ingress QoS policy uses dot1p, DEI with dot1p, or IP DSCP values to assign a forwarding class and profile to traffic, which facilitates the classification of traffic received on the access port. The user can assign a profile using DEI configured with a dot1p classification policy. The forwarding class is associated with meters or policers at ingress. The FC meters for unicast and multicast traffic and meter characteristics (for example, the CIR and PIR) are defined in the policy.

An access ingress QoS policy supports the definition of up to one (1) meter per forwarding class for unicast traffic, and up to one (1) meter per forwarding class for multipoint traffic (that is, broadcast, multicast, and unknown-unicast) for multipoint services. The definition of a maximum of 16 meters per access ingress QoS policy is supported.

For VPLS, the following four forwarding types (not to be confused with forwarding classes) are supported: unicast, multicast, broadcast, and unknown. The multicast, broadcast, and unknown traffic types are flooded to all destinations in the service and use the multipoint meter associated with the forwarding class. The unicast traffic type is handled in a point-to-point manner in the service and uses the unicast meter associated with the forwarding class.



Note:

An access ingress policy is supported only when the node is configured to operate in the high SAP scale mode using the **configure system resource-profile sap-scale-mode high** command on the 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE and **configure system global-res-profile sap-scale-mode high** command on the 7210 SAS-R6 and 7210 SAS-R12.

See the *7210 SAS-Mxp, S, Sx, T Services Guide* and *7210 SAS-R6, R12 Services Guide* for more information about high SAP scale mode.

See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information about the **sap-scale-mode** command.

20.1.1 Shared access ingress QoS policies on the 7210 SAS-Mxp

In addition to non-shareable access ingress port policies, users can configure and use shared access ingress policies. Shared policies optimize resource usage by associating a single instance of an access ingress QoS policy with multiple access ports. The software then allocates a single instance of classification entries and policing resources defined in the policy to classify and rate-limit ingress traffic received on all access ports. To use a shared access ingress QoS policy, enable the **config>system>resource-profile>qos-access-port-shared-res-mode** command. This mode is mutually exclusive to use non-shared mode. To use an access ingress policy exclusively with a port (to emulate the behavior of a non-shared access ingress policy), create a shared access ingress policy and associate it with only one port, instead of associating it with multiple ports.

With shared access ingress QoS policies, for traffic classification, the user has an option to use either table criteria or CAM criteria. With table criteria, the access ingress QoS policy uses dot1p, DEI with dot1p, or IP DSCP values to assign a forwarding class and profile to traffic. With CAM criteria, the access ingress QoS policy uses IP criteria to assign forwarding class to traffic. Shared access ingress policies provide similar capabilities for a policer definition as the non-shared access ingress policy.

20.1.2 Meter allocation rules for shared access ingress QoS policies

The following rules are used to allocate meters in a shared access ingress QoS policy.

- Sharing of a single meter for both unicast and multipoint traffic is allowed. Users can configure any of the available meters for multipoint traffic.
- When the **num-qos-classifiers** command is set to a value of 2, only a single meter is available for use for both unicast and multicast traffic. That is, a separate multipoint meter cannot be configured.
- When associating a meter with an FC for broadcast, unknown-unicast, or multicast (BUM) traffic, the system does not validate whether the meter is a multipoint meter. This allows a single meter to be used for unicast and BUM traffic, which allows for efficient use of access ingress QoS resources.
- The use of the **multipoint** keyword is not mandatory for meter creation. If a multipoint meter is specified during meter creation, Nokia recommends that it should be used only with BUM traffic when it is explicitly configured for an FC.
- Unlike the special semantics associated with meter 11 in a service ingress policy or meter 9 in non-shared access ingress QoS policy, there are no special semantics associated with meter 9 in a shared access ingress policy. In other words, meter 9 is not treated as the default multipoint meter for all FCs in a shared access ingress QoS policy. The user can configure meter 9 as a non-multipoint meter and use it for policing unicast traffic from any FC.

20.1.3 Resource allocation for non-shared access ingress QoS policies

This section describes the allocation of QoS resources from the CAM pool allocated to the policy to access ingress QoS policies using the **configure system resource-profile ingr-internal-tcam qos-access-port-ingress-resource** command.

For every FC in use, the system allocates two classification entries from the **qos-access-port-ingress-resource** CAM pool: one entry for the unicast meter and one for the multicast meter. Regardless of the number of dot1p entries in the dot1p classification policy and IP DSCP entries in the DSCP classification policy that map to a specific FC name, a maximum of two entries are required for the FC. That is, if the

user defines a dot1p classification policy with all seven dot1p values mapped to FC af, FC af needs only two classification entries in the CAM. An FC is considered to be in use if a dot1p classification entry, DSCP classification entry, or the default FC is mapped to the FC name.

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

The number of access ports that can be configured in **access-ingress-qos-mode port-mode** is limited by the number of classification resources available in the hardware and the number of access ports supported by the system.

20.1.3.1 Using index file for access ingress QoS policies

7210 SAS platforms use an index file to store the map that indicates the QoS resource allocation to access ports (for both non-shared and shared policies). After a reboot, the file is used to ensure appropriate resource allocation for all access ports that were using the access ingress policy before the reboot. In the absence of an index file, access ports that were configured successfully before the reboot may fail after the reboot. After the configuration file is saved, the index file is stored in flash memory. On system reboot, if the file is found, the system allocates resources as indicated in the stored map. If the file is not found, the system implements a best-fit algorithm and attempts to allocate resources for all access ports on a first-come-first-served basis. If the file is not found after reboot, the saved configuration may not execute successfully because resources may not be allocated to all access ports.



Note:

The index file used for the QoS map is different from the file used for storing interface indexes.

20.1.3.2 Calculating the number of QoS resources for non-shared access ingress QoS policies

To calculate the number of QoS resources used by a port-based access ingress QoS policy, the user must determine the number of FCs to use.

Only the FCs used by the match criteria classification entries configured in the dot1p and DSCP classification policies, which are referred to as "FCs in use," are considered to calculate the number of FCs.

Default unicast meter 1 and default multipoint meter 9, which are created by default when a new policy is created, cannot be deleted. Unless the user explicitly configures another unicast meter or multicast meter for the FCs, the default unicast meter 1 is used for all unicast traffic and default multipoint meter 9 is used for all multipoint traffic (that is, broadcast, multicast, and unknown-unicast).

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

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

- Two classification entries are used for the default *fc-name* configured using the **config qos access-ingress default-fc** command. The entries for the default FC are in addition to the FCs configured in the dot1p and DSCP classification policies.

Using the number of match criteria and FCs in use, calculate the total number of classification entries per policy with the following formula.

$$TC = \sum 2 * E(i) + 2$$

i = nc, h1, ef, h2, l1, af, l2, be

where:

- TC is the total number of classification entries per policy
- E(i) is the number of match criteria entries that classify packets to FCi. E(i) is one (1) if there is a dot1p or DSCP classification entry that classifies packets to FCi; otherwise, E(i) is zero.
- the number is multiplied by two for the number of classification entries that FCi requires
- another two entries are added for the default FC

20.1.3.3 Calculating the number of meters or policers for non-shared access ingress QoS policies

The total number of policers (TP) used is the number of meters configured in the policy. From the meters configured in the policy, only meters configured for use with an FC are counted for resource allocation. That is, meters that are created but not associated with an FC are not counted for resource allocation. A maximum of 16 meters are available per access ingress QoS policy.

20.1.3.4 Determining the number of resources allocated to non-shared access ingress QoS policies

The user must determine the value for the **config qos access-ingress num-qos-classifiers** command using the following formula.

$$\max (TC, TP)$$

See sections [Calculating the number of QoS resources for non-shared access ingress QoS policies](#) and [Calculating the number of meters or policers for non-shared access ingress QoS policies](#) for information about determining the values for TC and TP respectively.

20.1.3.5 Example of non-shared access ingress QoS policy resource calculations

The following sections describe example calculations for a non-shared access ingress QoS policy.

20.1.3.5.1 Example 1

The following output example shows the FC mapping in a configured dot1p classification policy, in which the FC is using the default meters (meter 1 for unicast and meter 9 for multicast).

Output example

```
configure> qos> dot1p-classification 10
dot1p 7 fc nc
dot1p 6 fc nc
default-dot1p-fc be profile out
exit
configure> qos> access-ingress 10
    dot1p-classification 10
    meter 1 create
        rate pir max cir 0
    exit
meter 9 multi-point create
    rate pir 100 cir 0
    exit
table-classification-criteria use-dot1p
    default-fc h1 profile in
exit
```

TC = 2 x 1 (for FC nc) + 2 x 1 (for **default-dot1p-fc** be) + 2 (for **default-fc** h1) = 6

TP = 2 x 2 = 4; (only 2 default meters)

The **num-qos-classifiers** value should be set to max (TC, TP) = max (6, 4) = 6

20.1.3.5.2 Example 2

The following output example shows the FC mapping in a configured dot1p classification policy, in which the FC is using the default meters (meter 1 for unicast and meter 9 for multicast).

Output example

```
configure> qos> dot1p-classification 10
dot1p 7 fc nc
dot1p 6 fc nc
default-dot1p-fc be profile out
exit
configure> qos> access-ingress 20
    dot1p-classification 10
    meter 1 create
        rate pir max cir 0
    exit
meter 9 multi-point create
    rate pir 100 cir 0
    exit
table-classification-criteria use-dot1p
    default-fc be profile out
exit
```

TC = 2 x 1 (for FC nc) + 2 x 1 (for **default-dot1p-fc** be) + 2 (for **default-fc** be) = 6

TP = 2 x 2 = 4; (only 2 default meters)

The **num-qos-classifiers** value should be set to max (TC, TP) = max (6, 4) = 6

20.1.3.5.3 Example 3

The following output example shows the configuration of the FC mapping in a dot1p classification policy, in which the FC is using both user-defined and default meters.



Note:

By default, unicast meter 1 and multipoint meter 9 are used if no explicit meter mapping is defined for the FC.

Example

```
configure> qos> dot1p-classification 10
dot1p 7 fc nc
dot1p 6 fc nc
default-dot1p-fc be profile out
exit

configure> qos> access-ingress 30
dot1p-classification 10
meter 1 create
rate pir max cir 0
exit
meter 9 multi-point create
rate pir 100 cir 0
exit
meter 2 create
rate 100 max cir 100
exit
meter 10 multi-point create
rate pir 100 cir 0
exit

fc nc
meter 2
multicast-meter 10
exit
table-classification-criteria use-dot1p
default-fc be profile out
exit
```

TC = 2 x 1 (for FC nc) + 2 x 1 (for **default-dot1p-fc** be) + 2 (for **default-fc** be) = 6

TP = 2 x 4 = 8; (2 default meters and 2 user-defined meters)

The **num-qos-classifiers** value should be set to max (TC, TP) = max (6, 8) = 8

20.1.4 Resource allocation for shared access ingress QoS policies on the 7210 SAS-Mxp

This section describes the allocation of QoS resources from the CAM pool allocated to the policy to shared access ingress QoS policies using the **config>system>resource-profile>ingr-internal-tcam>qos-access-port-shared-res** command.

The resource allocation varies based on the classification criteria configured for the policy. The following sections describe the resource allocation with **table-criteria** and **cam-criteria**.

With the **table-criteria** option, for every FC in use, the system allocates two classification entries from the **qos-access-port-shared-res** CAM pool: one entry to map the traffic for the unicast meter and one to map the traffic for the multicast meter.

Regardless of the number of dot1p entries in the dot1p classification policy and IP DSCP entries in the DSCP classification policy that map to a specific FC name, a maximum of two entries are required for the FC. That is, if the user defines a dot1p classification policy with all seven dot1p values mapped to FC af, FC af needs only two classification entries in the CAM. An FC is considered to be in use if a dot1p classification entry, DSCP classification entry, or the default FC is mapped to the FC name. For every meter or policer in use, the system allocates one meter in the hardware. A meter or policer is considered to be in use when it is associated with an FC in use.

With the **cam-criteria** option, for every classification entry and FC in use, the system allocates two classification entries from the **qos-access-port-shared-res** CAM pool; one entry for the unicast meter and one for the multicast meter. That is, if the user defines a IP criteria with 10 entries, then all the classification entries are configured to FC "af", FC "af" needs two classification entries in the CAM per each IP criteria entry for a total of 20 entries in the CAM. An FC is considered to be in use if a **cam-criteria** classification entry or the default FC is mapped to the FC name. For every meter or policer in use, the system allocates one meter in the hardware. A meter or policer is considered to be in use when it is associated with an FC in use.

The number of access ports that can be configured in **access-ingress-qos-mode port-mode** is limited by the number of classification resources available in the hardware and the number of access ports supported by the system.

20.1.4.1 Calculating the number of QoS resources for shared access ingress QoS policies

To calculate the number of QoS resources used by a port-based shared access ingress QoS policy, the user must determine the number of FCs to use and the criteria used for classification of traffic to FC.

If using **table-criteria**, the QoS resource calculation is exactly the same as the non-shared access ingress QoS policy described in [Calculating the number of QoS resources for non-shared access ingress QoS policies](#). Resources are shared by all access ports the access ingress QoS policy is associated with.

If using **cam-criteria**, the QoS resource calculation for a port-based shared access ingress QoS policy using **cam-criteria**, the user must determine the number of FCs and the number of classification entries to use for classification of traffic to the FC.

Only the FCs used by the match criteria classification entries configured in the policy, which are referred to as "FCs in use," are considered in the calculation of the number of FCs.

The default unicast meter 1, which is created by default when a new policy is created, cannot be deleted. Unless the user explicitly configures another unicast meter or multicast meter for the FCs, the default unicast meter 1 is used for all unicast and multipoint traffic (for example, broadcast, multicast, and unknown-unicast).

The following meter selection rules are used by different traffic types under various configurations to compute the number of classification entries per FC in use:

- Only meter 1 is defined in the default policy; Meter 9 is not created. All FC and all traffic types use meter 1 by default. The following is a sample configuration output:

```
*7210-SAS>config>qos# access-ingress share-resources 65535 create
*7210-SAS>config>qos>access-ingress$ info
-----
    num-qos-classifiers 2
    meter 1 create
    exit
-----
*7210-SAS>config>qos>access-ingress$
```

- If an FC is created without explicit meters, the default meter 1 is used for unicast traffic and BUM traffic types (for example, broadcast, multicast, and unknown-unicast traffic). When **cam-criteria** is configured, this uses a single entry in the CAM for each classification entry configured in the policy. When **table-criteria** is configured, this uses a single entry in the CAM for each FC.
- If an FC is created with only an explicit unicast meter, that meter is used for unicast traffic and for BUM traffic types. When **cam-criteria** is configured, this uses a single entry in the CAM for each classification entry configured in the policy. When **table-criteria** is configured, this uses a single entry in the CAM for each FC.
- If an FC is created with an explicit unicast meter and explicit multicast meter, use the configured unicast meter for unicast traffic and the configured multicast meter for all other traffic types. When **cam-criteria** is configured, this uses two entries in the CAM for each classification entry configured in the policy. When **table-criteria** is configured, this uses two entries in the CAM for each FC.
- If an FC is created with only an explicit multicast meter, use the default meter 1 for unicast traffic and the configured multicast meter for BUM traffic. When **cam-criteria** is configured, this uses two entries in the CAM for each classification entry configured in the policy. When **table-criteria** is configured, this uses two entries in the CAM for each FC.
- If an FC is configured explicitly to use a unicast meter and a multicast meter and both the traffic types are configured to use the same meter ID, the software allocates two entries in the CAM. When **cam-criteria** is configured, this uses two entries in the CAM for each classification entry configured in the policy. When **table-criteria** is configured, this uses two entries in the CAM for each FC.

Using the number of match criteria and FCs in use, calculate the total number of classification entries per policy with the following formula.

$$TC = \sum (M(i) * E(i) + 2)$$

i = 1 – maximum number of classification entries

where:

- TC is the total number of classification entries per policy
- E(i) is the number of match criteria entries that classify packets to FCi. The E(i) is one (1) if there is a CAM criteria classification entry that classifies packets to FCi; otherwise, the E(i) is zero.
- The M(i) is the number of meters FCi requires and determined using the meter selection rules provided in the preceding list. Each meter defined for the FC consumes a CAM entry to map the traffic flow to defined FCi.
- two additional entries are added for the default FC

20.1.4.2 Calculating the number of meters or policers for shared access ingress QoS policies

The total number of policers (TP) used is the number of meters configured in the policy. Only meters configured for use with an FC are counted for resource allocation. That is, meters that are created but not associated with an FC are not counted. A maximum of 16 meters are available per access ingress QoS policy.

20.1.4.3 Determining the number of resources allocated to shared access ingress QoS policies

The user must determine the value for the **config>qos>access-ingress>num-qos-classifiers** command using the following formula.

$$\max(\text{TC}, \text{TP})$$

where:

- TC is the total number of classification entries per policy
- TP is the total number of policers

See sections [Calculating the number of QoS resources for shared access ingress QoS policies](#) and [Calculating the number of meters or policers for shared access ingress QoS policies](#) for information about determining the values for the TC and TP parameters.

20.1.4.4 Examples of shared access ingress QoS policy resource calculations

This section provides example resource calculations for a shared access ingress QoS policy.

20.1.4.4.1 Example 1

The following output is an example of the FC mapping for a **cam-criteria** (IP criteria) access ingress QoS policy, in which the FC is using the the default meters (meter 1).

Example

```
configure> qos> access-ingress 10
meter 1 create
rate pir max cir 0
exit
classification-criteria cam-criteria
fc h1
  meter 1
exit
fc nc
  meter 1
exit
default-fc h1 profile in
ip-criteria
  entry 1
    action fc nc
    match src-ip 192.168.1.1/32
  exit
exit
```

where:

$$\text{TC} = 1 \times 1 \text{ (for FC nc)} + 1 \times 1 \text{ (for default-fc h1)} = 2$$

$$\text{TP} = 1 \times 2 = 2; \text{ (only 1 default meter)}$$

The **num-qos-classifiers** value should be set to $\max(\text{TC}, \text{TP}) = \max(2, 2) = 2$.

20.1.4.4.2 Example 2

The following output is an example of the FC mapping for a **cam-criteria** (IP criteria) access ingress QoS policy, in which the FC is using the non-default meters (meter 1 and meter 3).

Example

```
configure> qos> access-ingress 10
meter 1 create
rate pir max cir 0
exit
meter 3 create
rate pir max cir 10
exit
fc h1
  meter 1
  multicast-meter 3
exit
fc nc
  meter 1
exit

classification-criteria cam-criteria
default-fc h1 profile in
ip-criteria
  entry 1
    action fc nc
    match src-ip 192.168.1.1/32
  exit
exit
```

where:

TC = 1 x 1 (for FC nc) + 1 x 2 (for **default-fc** h1) = 3

TP = 2 x 2 = 4; (2 meters; meter 3 and meter 1)

The **num-qos-classifiers** value should be set to max (TC, TP) = max (3, 4) = 4.

20.1.4.4.3 Example 3

The following output is an example of the FC mapping for a **cam-criteria** (IP criteria) access ingress QoS policy in which a mix of FCs are in use and where some FCs use the non-default meters (meter 1 and meter 3) and some use the default meter 1.



Note: By default, unicast meter 1 is used if no explicit meter mapping is defined for the FC.

Example

```
configure> qos> access-ingress 10
meter 1 create
rate pir max cir 0
exit
meter 3 create
rate pir max cir 10
exit
fc h1
```



```

meter 1
  multicast-meter 3
exit
fc nc
  multicast-meter 9
exit
fc af
  multicast-meter 4
exit
fc be
exit

classification-criteria cam-criteria
default-fc be profile in
ip-criteria
  entry 1
    action fc nc
    match src-ip 192.168.1.1/32
  exit
  entry 10
    action fc af
    match dst-ip 192.168.100.1/32
  exit
  entry 20
    action fc h1
    match src-ip 192.168.200.1/32
  exit
  entry 30
    action fc af
    match dst-ip 192.168.300.1/32
  exit
  entry 40
    action fc af
    match src-ip 192.168.400.1/32
  exit
exit

```

where:

TC = 1 x 2 (for FC nc) + 3 x 2 (for FC af) + 1 x 1 (for **default-fc** be) = 11

TP = 2 x 2 = 4; (2 meters; meter 3 and meter 1)

The **num-qos-classifiers** value should be set to max (TC, TP) = max (11, 4) = 11 (rounded off to the nearest even number greater than 11, which is 12).

20.1.5 Configuration guidelines for a port-based access ingress QoS policy

Use the following guidelines to configure a port-based access ingress QoS policy.



Note: The following guidelines are generic and some of the commands are not available on all platforms.

- Delete all SAPs on the port or the LAG before switching between the **sap-mode** and **port-mode** options for the **access-ingress-qos-mode** command.
- Perform the following steps before changing the **access-ingress-qos-mode** command to **port-mode**:
 1. Reboot the node in **sap-scale-mode high**.

2. If shared mode is enabled (**config>system>resource-profile>qos-access-port-shared-res-mode**), ensure that the following resources are allocated:
 - **config>system>resource-profile>ingress-internal-tcam>qos-access-ingress-resource**
 - **config>system>resource-profile>ingress-internal-tcam>qos-access-port-shared-res**
 3. Configure an access port using the **config>port>ethernet>mode** command with the **access** option specified.
- If the **config>port>ethernet>access-ingress-qos-mode** command is set to **port-mode**, access ingress policy 1 is attached to the port by default. In shared mode, the access ingress policy ID 65536 is attached to the port by default. The policy can be replaced with a user-defined access ingress QoS policy provided that sufficient resources are available in the **qos-access-port-ingress-resource** slice or **qos-access-port-shared-res** slice in shared mode.
 - Resource allocation using **qos-access-port-ingress-resource** has no restrictions and resources can be configured using this command in **sap-scale-mode low**, **high**, or in shared mode; however, allocating resources in **sap-scale-mode low** or in shared mode wastes resources because an access ingress QoS policy can only be attached in **sap-scale-mode high** and resources for shared policies are allocated from the **qos-access-port-shared-res** slice.
 - Reset the user-defined access ingress QoS policy to 1 or 65536 on the port before changing the **access-ingress-qos-mode** command to **sap-mode**.
 - If the **access-ingress-qos-mode** command is configured in the **config>lag** context, reset the user-defined access ingress QoS policy to 1 or 65536 on the primary LAG member before changing the **access-ingress-qos-mode** command to **sap-mode**.

20.1.6 Basic configurations for non-shared access ingress QoS policies

A basic non-shared access ingress QoS policy must conform to the following configuration rules:

- The policy must have a unique access ingress QoS policy ID.
- The policy can specify meter parameters, such as CIR and PIR, to define meters for use.
- The policy may include a dot1p and DSCP template attachment to map dot1p and IP DSCP values to the FC.
- Optionally, the policy may configure either the IP DSCP or dot1p or both for classification. The user can also assign an ingress profile based on either a dot1p, DEI with dot1p, or IP DSCP.

Output example

The following output is an example configuration of an access ingress QoS policy.

```
*A:Dut-A>config>qos# access-ingress 10
*A:Dut-A>config>qos>access-ingress# info
-----
meter 1 create
  mode trtcm2
  adaptation-rule cir min
  rate cir 1010 pir 4040
  color-mode color-blind
  mbs 512 bytes
  cbs 100 kbytes
exit
meter 3 create
exit
```

```
meter 5 multi-point create
exit
meter 9 multi-point create
exit
fc "ef" create
    meter 3
    multicast-meter 5
exit
counter-mode forward-drop-count
default-fc "ef" profile in
dot1p-classification 40
dscp-classification 30
table-classification-criteria use-dscp
num-qos-classifiers 8
-----
```

20.1.6.1 Editing a non-shared access ingress QoS policy configuration

The user can edit existing policies and entries through the CLI or NMS. The changes are applied immediately to all services to which the policy applies.



Note:

The *num-resources* parameter value for the **config>qos>access-ingress>num-qos-classifiers** command cannot be modified when the policy is in use.

Perform the following to prevent configuration errors:

- Copy the policy to a work area.
- Edit the policy.
- Overwrite the original policy.

20.1.6.2 Removing a non-shared policy from the QoS configuration

Use the following syntax to remove an access ingress policy from the QoS configuration.

```
config>qos# no access-ingress policy-id
```

Example:

```
config>qos# no access-ingress 100
config>qos# no access-ingress 1010
```

20.1.6.3 Deleting a non-shared access ingress QoS policy

Every access Ethernet port is associated, by default, with the default access ingress policy (*policy-id* 1) when the **access-ingress-qos mode** command is set to **port-mode**. You can replace the default policy with a user-defined policy. If the non-default access ingress policy is removed, the association reverts to default policy-id 1. A QoS policy cannot be deleted until it is removed from all access ports where it is applied.

```
*A:card-1>config>qos# no access-ingress 30
```

MINOR: CLI Could not remove Access ingress policy "30" because it is in use.



Note:

The **access-ingress-qos-mode** command can only be changed to **sap-mode** if access ingress policy 1 is attached to the port.

20.1.7 Basic configurations for shared access ingress QoS policies

A basic shared access ingress QoS policy must conform to the following configurations rules:

- The policy must have a unique access ingress QoS policy ID.
- The policy can specify meter parameters, such as CIR and PIR, to define meters for use.
- The policy may include a dot1p and DSCP template attachment to map dot1p and IP DSCP values to the FC.
- Optionally, the policy may configure either the IP DSCP or dot1p or both for classification. The user can also assign an ingress profile based on either a dot1p, DEI with dot1p, or IP DSCP.

Output example

The following output is an example configuration of a shared access ingress QoS policy.

```
*A:Mxp> configure qos access-ingress 65536
*A:Mxp>config>qos>access-ingress# info detail
-----
      description "Default ACCESS ingress shared QoS policy."
      meter 1 create
        mode trtcml
        adaptation-rule cir closest pir closest
        rate cir 0 pir max
        color-mode color-aware
        mbs default kbits
        cbs default kbits
      exit
      classification-criteria table-criteria
      counter-mode in-out-profile-count
      default-fc "be" profile out
      dot1p-classification 1
      dscp-classification 1
      table-classification-criteria both-dscp-dot1p
      num-qos-classifiers 2
-----
```

20.1.7.1 Editing a shared access ingress QoS policy configuration

The user can edit existing policies and entries through the CLI or NMS. The changes are applied immediately to all services to which the policy applies.



Note:

The *num-resources* parameter value for the **config>qos>access-ingress>num-qos-classifiers** command cannot be modified when the policy is in use.

Perform the following to prevent configuration errors:

- Copy the policy to a work area.

- Edit the policy.
- Overwrite the original policy.

20.1.7.2 Removing a shared policy from the QoS configuration

Use the following syntax to remove an access ingress policy from the QoS configuration.

```
config>qos# no access-ingress policy-id
```

Example:

```
config>qos# no access-ingress 100
config>qos# no access-ingress 1010
```

20.1.7.3 Deleting a shared access ingress QoS policy

Every access Ethernet port is associated, by default, with the default access ingress shared QoS policy (policy ID 65536) when the **access-ingress-qos-mode** command is set to **port-mode**. Users can replace the default policy with a user-defined policy. If the non-default access ingress shared QoS policy is removed, the association reverts to the default policy ID 65536. A QoS policy cannot be deleted until it is removed from all access ports where it is applied.

```
Mxp>config>qos# no access-ingress 20
MINOR: QOS #11031 Access ingress policy is already in use
```



Note:

The **access-ingress-qos-mode** command can only be changed to **sap-mode** if access ingress policy 65535 is attached to the port.

20.2 Access-ingress QoS policy command reference

20.2.1 Command hierarchies

- [Access-ingress QoS configuration commands](#)
- [Show commands](#)

20.2.1.1 Access-ingress QoS configuration commands

```
config
- qos
  - access-ingress policy-id [share-resources] [create]
  - no access-ingress policy-id
    - classification-criteria classification-criteria
    - no classification-criteria
    - counter-mode {in-out-profile-count | forward-drop-count}
    - no counter-mode
```

```
- description description-string
- no description
- default-fc fc-name profile {in | out}
- no default-fc
- dot1p-classification description-string
- no dot1p-classification
- dscp-classification description-string
- no dscp-classification
- fc fc-name [create]
- no fc fc-name
  - meter meter-id
  - no meter
  - multicast-meter meter-id
  - no multicast-meter
- ip-criteria [use-port-range]
- no ip-criteria
  - entry entry-id [create]
  - no entry entry-id
    - action [fc fc]
    - no action
    - description description-string
    - no description
    - match [protocol protocol-id]
    - no match
      - dscp dscp-value | dscp-name [dscp-mask]
      - no dscp
      - dst-ip ip-address [ipv4-address-mask]
      - dst-ip ip-address/mask
      - no dst-ip
      - dst-port {eq} dst-port-number
      - dst-port {range} dst-port-number1 dst-port-number2
      - no dst-port
      - ip-prec ip-prec-value [ip-prec-mask]
      - no ip-prec
      - src-ip ip-address/mask
      - src-ip ip-address [ipv4-address-mask]
      - no src-ip
      - src-port {eq} src-port-number
      - src-port {range} src-port-number1 src-port-number2
      - no src-port
    - renum old-entry-id new-entry-id
  - meter meter-id [create] [multi-point]
- no meter meter-id
  - adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
  - no adaptation-rule
  - cbs size [kbits | bytes | kbytes]
  - no cbs
  - color-mode {color-aware | color-blind}
  - no color-mode
  - mbs size [kbits | bytes | kbytes]
  - no mbs
  - mode mode
  - no mode
  - rate [cir cir-rate] [pir pir-rate]
  - no rate
- num-qos-classifiers num-resources
- no num-qos-classifiers
- table-classification-criteria table-classification-criteria
```

20.2.1.2 Show commands

```
show
- qos
  - access-ingress [policy-id] association
  - access-ingress [policy-id] [detail]
```

20.2.2 Command descriptions

20.2.2.1 Generic commands

access-ingress

Syntax

access-ingress *policy-id* [**share-resources**] [**create**]
no access-ingress *policy-id*

Context

config>qos

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command creates or edits an access ingress QoS policy, which is used when the **access-ingress-qos-mode** command is set to **port-mode**. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for more information about the **access-ingress-qos-mode** command.

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

Default

access-ingress 1 (for non-shared mode)
access-ingress 65536 (for shared mode)

Parameters

policy-id
Specifies the access-ingress policy ID.

Values	1 to 65535
	1 to 65536 (only on the 7210 SAS-Mxp)



Note: The value of 65536 is the default policy reserved for use when the node is configured to share access-ingress QoS policies across ports and is not modifiable by the user in the description.

create

Keyword to create an access-ingress policy. The create keyword requirement can be enabled or disabled in the **environment>create** context.

share-resources

Keyword to indicate that if the policy is associated with a port, it shares existing resources if another port has already been instantiated with this policy. If this keyword is not specified, the port allocates new resources if they are available and fails if there are no resources. This keyword is only available on the 7210 SAS-Mxp and can only be used when the **config>system>resource-profile>qos-access-port-shared-res-mode** command is enabled.

description

Syntax

description *description-string*

no description

Context

config>qos>access-ingress

config>qos>access-ingress>ip-criteria>entry (only on the 7210 SAS-Mxp)

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command creates a text description stored in the configuration file for a configuration context. The text string is stored in the configuration file and identifies the context in the file.

The **no** form of this command removes the description string for the configuration context.

Parameters

description-string

Specifies a text string that describes 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.

20.2.2.2 Access ingress QoS policy commands

classification-criteria

Syntax

classification-criteria *classification-criteria*
no classification-criteria

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp

Description

This command specifies whether table-based or CAM-based classification criteria is used for the access ingress QoS policy. The *classification-criteria* value cannot be changed when the policy is associated with a port.

Dot1p classification and IP DSCP classification policy-based classification entries can be used when using table-based classification criteria.

IP criteria classification entries can be used when using CAM-based classification criteria.

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

Default

classification-criteria table-criteria

Parameters

classification-criteria

Specifies if table-based or CAM-based classification criteria is used for the access ingress QoS policy.

Values cam-criteria, table-criteria

counter-mode

Syntax

counter-mode {in-out-profile-count | forward-drop-count}
no counter-mode

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command allows the user to set the counter mode for the counters associated with access port ingress meters (also known as policers). A pair of counters is available with each meter. These counters count different events based on the counter mode value.

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

Default

counter-mode in-out-profile-count

Parameters

in-out-profile-count

Specifies that the in-profile and out-profile packets and octets are counted per meter.

fwd-drop-count

Specifies that the forwarded and dropped packets and octets are counted per meter.

default-fc

Syntax

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

no default-fc

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command defines or edits the default forwarding class to be assigned to packets that do not match the explicitly configured classification entries. See the [table-classification-criteria](#) command for more information about configuring the default FC.

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

Default

default-fc be profile out

Parameters

fc-name

Specifies the forwarding class name.

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

profile {in | out}

Specifies whether packets are in-profile or out-of-profile. All packets that are assigned to this forwarding class are considered in or out of profile based on which keyword is configured. A value of **in** defines the packet as being in-profile and a value of **out** defines the packet as being out-of-profile.

dot1p-classification

Syntax

dot1p-classification *policy-id*

no dot1p-classification

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command associates a dot1p classification policy with an access ingress QoS policy. See [IP DSCP and dot1p classification policy support](#) for more information about dot1p classification policies.

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

Default

dot1p-classification 1

Parameters

policy-id

Specifies the dot1p classification policy ID.

Values 1 to 65535

dscp-classification

Syntax

dscp-classification *policy-id*

no dscp-classification

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command associates a DSCP classification policy with an access ingress QoS policy. See [IP DSCP and dot1p classification policy support](#) for more information about dot1p classification policies.

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

Default

dscp-classification 1

Parameters

policy-id

Specifies the DSCP classification policy ID.

Values 1 to 65535

fc

Syntax

fc *fc-name* [create]

no fc *fc-name*

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command creates a class instance of the forwarding class. After the *fc-name* is created, classification actions can be applied to the forwarding class and it can be used in classification criteria configured in the DSCP and dot1p classification policies.

The **no** form of this command removes all explicit meter mappings for *fc-name* forwarding types. The meter mappings revert to the default meters for *fc-name*.

Default

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

Parameters

fc-name

Specifies the case-sensitive, system-defined FC name for which policy entries will be created.

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

create

Keyword to create the forwarding class. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

meter

Syntax

meter *meter-id*

no meter

Context

config>qos>access-ingress>fc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command overrides the default unicast forwarding type meter mapping for **fc** *fc-name*. The specified meter ID must exist within the policy as a non-multipoint meter before the mapping can be made. After the FC mapping is applied, the meter ID is used to forward all unicast traffic that uses this policy on the port.

The **no** form of this command reverts the unicast (point-to-point) meter ID to the default meter for the FC.

Default

meter 1

Parameters

meter-id

Specifies the meter ID, which must be an existing, non-multipoint meter defined in the **config>qos>access-ingress** context.

Values 1 to 16

multicast-meter

Syntax

multicast-meter *meter-id*

no multicast-meter

Context

config>qos>access-ingress>fc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command overrides the default multicast forwarding type meter mapping for **fc** *fc-name*. The specified meter ID must exist within the policy as a multipoint meter before the mapping can be made. After the FC mapping is applied, the meter ID is used to forward all multicast traffic that uses this policy on the port.

The **no** form of this command reverts the multicast forwarding type *meter-id* to the default meter for the FC.

Default

multicast-meter 9

Parameters

meter-id

Specifies the multicast meter, which must be an existing, multipoint meter defined in the **config>qos>access-ingress** context.

Values 1 to 16

ip-criteria

Syntax

ip-criteria [**use-port-range**]

no ip-criteria

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp

Description

This command creates or edits policy entries that specify an IP criteria (such as, IP DSCP, IP addresses, and so on.). IP criteria-based access ingress QoS policies are used to select the appropriate ingress meter and corresponding forwarding class for matched traffic.

The 7210 SAS implementation exits on the first match found and executes 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 entries specified under this node. Once IP criteria entries are removed from an access ingress QoS policy, the IP criteria is removed from all services where that policy is applied.

Parameters

use-port-range

Specifies that classification entries are used to specify range values for TCP/UDP source and destination port fields, in addition to other IP criteria fields. If an IP criteria classification policy is created without this option, the filter entries can only use exact values for TCP/UDP source and destination port fields. By default, an IP criteria classification policy is created without the **use-port-range** option.

There are limited amount of unique port range values that can be used (see the **tools>dump>system-resources** command to know the current usage). The port range entries are shared among QoS and ACLs, along with source and destination port range values.

entry

Syntax

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

no entry *entry-id*

Context

config>qos>access-ingress>ip-criteria

Platforms

7210 SAS-Mxp

Description

This command creates or edits an IP criteria entry for the policy. Multiple entries can be created using unique *entry-id* numbers.

The list of flow criteria is evaluated in a top-down fashion with the lowest entry ID at the top and the highest entry ID at the bottom. If the defined match criteria for an entry within the list matches the information in the egress packet, the system stops matching the packet against the list and performs the reclassification actions for the matching entries. If none of the entries match the packet, the IP flow reclassification list has no effect on the packet.

An entry is not populated in the list unless the **action** command is executed for the entry. An entry that is not populated in the list has no effect on egress packets. If the **action** command is executed without any explicit reclassification actions specified, the entry is populated in the list allowing packets matching the entry to exit the list, preventing them from matching entries lower in the list. Because 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 is not reclassified.

The **no** form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services where that policy is applied.

Parameters

entry-id

Specifies a match criterion and the corresponding action. It is recommended that multiple entries be given an *entry-id* in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword **action fc fc-name** for it to be considered complete. Entries without the **action** keyword are considered incomplete and are rendered inactive.

Values 1 to 250

create

Keyword to create a flow entry when the system is configured. 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

Syntax

action [**fc** *fc*]

no action

Context

config>qos>access-ingress>ip-criteria>entry

Platforms

7210 SAS-Mxp

Description

This mandatory command associates the forwarding class with a specific IP criteria entry ID. The **action** command supports setting the forwarding class parameter. Packets that meet all match criteria within the entry have their forwarding class overridden based on the parameters included in the **action** parameters.

The **action** command must be executed for the match criteria to be added to the active list of entries.

Each time action is executed on a specific entry ID, the previous entered values for **fc** are overridden with the newly defined parameters.

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 are lost.

Default

action specified by the **default-fc**

Parameters

fc

Specifies the forwarding class name for the queue or meter. The value given for *fc* must be one of the predefined forwarding classes in the system.

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

match

Syntax

match [**protocol** *protocol-id*]

Context

config>qos>access-ingress>ip-criteria>entry

Platforms

7210 SAS-Mxp

Description

This command configures match criteria for access ingress QoS policies. When the match criteria have been satisfied, the action associated with the match criteria is executed.

Parameters

protocol-id

Specifies an IP protocol used as an access ingress QoS policy match criterion.

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

Values 0 to 255

dscp

Syntax

dscp *dscp-value* | *dscp-name* [*dscp-mask*]

no dscp

Context

config>qos>access-ingress>ip-criteria>entry>match

Platforms

7210 SAS-Mxp

Description

This command configures a DiffServ Code Point (DSCP) to be used for classification of packets from the specified FC.

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



Note:

- Either the DSCP name or DSCP value with a mask can be configured.
- When the user configures the *dscp* value alone, the **show** command displays the *dscp* value as the configured value and the *dscp-mask* as the FC.

Parameters

dscp-value

Specifies the DSCP value in hexadecimal, decimal, or binary format.

Values 0 to 63

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, 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

dscp-mask

Specifies a 6-bit mask that 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 4 and 0b000100 for value and mask.

Values 0 to 63 (decimal, hexadecimal, or binary)

Default 63 (exact match)

dst-ip

Syntax

dst-ip *ip-address* [*ipv4-address-mask*]

dst-ip *ip-address/mask*

no dst-ip

Context

config>qos>access-ingress>ip-criteria>entry>match

Platforms

7210 SAS-Mxp

Description

This command configures a destination address range to be used as an access ingress QoS policy match criterion.

To match on the destination address, specify the address and its associated mask; for example, 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

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

Parameters

ip-address

Specifies the IPv4 address of the destination IP. This address must be unique within the subnet and specified in dotted-decimal notation.

Values 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)

dst-port

Syntax

no dst-port

dst-port {*eq*} *dst-port-number*

dst-port {*range*} *dst-port-number1* *dst-port-number2*

Context

config>qos>access-ingress>ip-criteria>entry>match

Platforms

7210 SAS-Mxp

Description

This command configures a destination TCP or UDP port number or a range of port numbers for an access ingress QoS policy match criterion.



Note:

- The IP criteria must be created with the **use-port-range** option to configure a range of port numbers.
- A limited number of unique port range values are available for use (see the **tools>dump>system-resources** command for the current usage). The port range entries are shared among the QoS classification and ACL entries, and among source and destination port range values. The available entries in the hardware port range resource pool is allocated on a first come, first served basis. The following rules apply to the allocation of these port range pool entries:
 - One port range hardware entry is required to match a unique combination of port range values configured by the user. Two port range values (for example, port range 1024 to 2048 and port range 3050 to 3055) are considered to be unique if their range 1 and range 2 values do not match. Each unique port range value consumes one entry each in hardware.
 - Each unique port range entry requires two entries in the hardware table if it is used for both a source and destination match. In other words, a unique port range value requires one entry for matching a source port and another entry for matching a destination port.
 - If one or more QoS or ACL policy use the same unique port range value for either the source port match or destination port match but not both, only a single entry in the hardware table is required. For example, if port range 1000 to 2000 is used in both the access ingress policy and filter policy to match on source port, it requires one entry in the hardware port range table.
 - If one or more QoS and or ACL policy use the same unique port range value for both source port match and destination port, two entries in the hardware table are required. For example, if port range 1000 to 2000 is used in both access ingress policy and filter policy to match on both the source port and destination port, it requires two entries in the hardware port range table.

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

Parameters

eq *dst-port-number*

Specifies the TCP or UDP port number to match, specified as equal to (**eq**) the destination port value specified as a decimal integer.

Values 1 to 65535 (decimal hex or binary)

range *dst-port-number1*

Specifies the first destination port number in the range to match a range of ports.

Values 1 to 65535 (decimal hex or binary)

range *dst-port-number2*

Specifies the last destination port number in the range to match a range of ports. The second number must be greater than the first number in the range.

Values 1 to 65535 (decimal hex or binary)

ip-prec

Syntax

ip-prec *ip-prec-value* [*ip-prec-mask*]

no ip-prec

Context

config>qos>access-ingress>ip-criteria>entry>match

Platforms

7210 SAS-Mxp

Description

This command defines a specific IP precedence value that must be matched to perform the associated classification actions. If an ingress packet on the port where the access ingress QoS policy is applied to matches the specified IP precedence value, the actions associated with this entry are taken.

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 QoS behavior. The precedence bits are also part of the newer 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. A match entry cannot match on both IP DSCP and IP precedence values. That is, the user can use either an IP DSCP or IP precedence match in a match entry, but not both. The software blocks configuration of an IP precedence match if **ip-dscp** is configured already. The converse is also true. A single policy having multiple match entries can have entries so that some of them match IP DSCP and some others match IP precedence. The order of the entry determines the priority of the match.

The **no** form of this command removes the IP Precedence match criterion.

Parameters

ip-prec-value

Specifies the unique IP header ToS byte precedence bits value that match the IP precedence rule.

Values 0 to 7

ip-prec-mask
Specifies the mask to use for the match.

Values 0 to 7

src-ip

Syntax

src-ip *ip-address/mask*
src-ip *ip-address [ipv4-address-mask]*
no src-ip

Context

config>qos>access-ingress>ip-criteria>entry>match

Platforms

7210 SAS-Mxp

Description

This command configures a source IPv4 address range to be used as an access ingress QoS policy match criterion.

To match on the source IPv4 address, specify the address and its associated mask; for example, 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

The **no** form of this command removes the source IPv4 address match criterion.

Parameters

ip-address
Specifies the IPv4 address of the destination IP. This address must be unique within the subnet and specified in dotted-decimal notation.

Values 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)

src-port

Syntax

src-port {eq} *src-port-number*

src-port {range} *src-port-number1 src-port-number2*

no src-port

Context

config>qos>access-ingress>ip-criteria>entry>match

Platforms

7210 SAS-Mxp

Description

This command configures a source TCP or UDP port number or a range of port numbers for an access ingress QoS policy match criterion.



Note:

- The IP criteria must be created with the **use-port-range** option to configure a range of port numbers.
- A limited number of unique port range values are available for use (see the **tools>dump>system-resources** command for the current usage). The port range entries are shared among the QoS classification and ACL entries, and among source and destination port range values. The available entries in the hardware port range resource pool is allocated on a first come, first served basis. The following rules apply to the allocation of these port range pool entries:
 - One port range hardware entry is required to match a unique combination of port range values configured by the user. Two port range values (for example, port range 1024 to 2048 and port range 3050 to 3055) are considered to be unique if their range 1 and range 2 values do not match. Each unique port range value consumes one entry each in hardware.
 - Each unique port range entry requires two entries in the hardware table if it is used for both a source and destination match. In other words, a unique port range value requires one entry for matching a source port and another entry for matching a destination port.
 - If one or more QoS or ACL policy use the same unique port range value for either the source port match or destination port match but not both, only a single entry in the hardware table is required. For example, if port range 1000 to 2000 is used in both the access ingress policy and filter policy to match on source port, it requires one entry in the hardware port range table.
 - If one or more QoS and or ACL policy use the same unique port range value for both source port match and destination port, two entries in the hardware table are required. For example, if port range 1000 to 2000 is used in both access ingress policy and filter policy to

match on both the source port and destination port, it requires two entries in the hardware port range table.

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

Parameters

eq *src-port-number*

Specifies the TCP or UDP port number to match, specified as equal to (**eq**) the destination port value specified as a decimal integer.

Values 1 to 65535 (decimal hex or binary)

range *src-port-number1*

Specifies the first source port number in the range to match a range of ports.

Values 1 to 65535 (decimal hex or binary)

range *src-port-number2*

Specifies the last source port number in the range to match a range of ports. This value must be greater than the first value specified in the range command.

Values 1 to 65535 (decimal hex or binary)

renum

Syntax

renum *old-entry-id new-entry-id*

Context

config>qos>access-ingress>ip-criteria

Platforms

7210 SAS-Mxp

Description

This command renumbers existing QoS policy criteria entries to properly sequence policy entries.

This can be required in some cases because the 7210 SAS exits when the first match is found and executes the actions in accordance with the accompanying **action** command. This requires that entries be sequenced correctly from most to least explicit.

Parameters

old-entry-id

Specifies the entry ID of the existing QoS policy criteria entry.

Values 1 to 250

new-entry-id

Specifies the entry ID for the new QoS policy criteria entry.

Values 1 to 250

meter

Syntax

meter *meter-id* [**create**] [**multi-point**]

no meter *meter-id*

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command configures an access ingress QoS policy meter (also known as a policer). The **meter** command allows the creation of unicast and multipoint meters.

A meter can be shared by multiple FCs, but the unicast and multicast traffic of an FC cannot share the same meter. That is, two or more FCs can share the same unicast and multicast meter for unicast and multicast traffic, but a minimum of 2 meters are required; one for unicast traffic and another for multicast traffic.

Multipoint meters receive ingress packets destined for multiple destinations, and handle traffic bound to these destinations. In non-multipoint services, such as Epipe services, all traffic is considered unicast because of the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service is not mapped to a multipoint service meter.

The **no** form of this command removes the meter ID from the access ingress QoS policy and from existing ports that use the policy. If any forwarding class forwarding types are mapped to the meter, they revert to their default meters. When a meter is removed, pending accounting information for each port meter created when the meter is defined in the policy is discarded.

Default

meter 1 (for unicast traffic)

meter 9 multipoint (for traffic other than unicast traffic)

Parameters

meter-id

Specifies the meter ID that uniquely identifies the meter within the policy. This is a required parameter and must be specified each time the **meter** command is run.

Values 1 to 16

multipoint

Specifies that the defined *meter-id* is for multipoint forwarded traffic only. This *meter-id* can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. Attempting to map forwarding class unicast traffic to a multipoint meter causes the system to generate an error and the current unicast traffic meter mapping is unchanged.

A meter must be defined as multipoint when it is created using the **create** keyword. Applying the keyword after the meter is created is not allowed. Attempting to modify the command to include the **multipoint** keyword will cause the system to generate an error and the command will not execute.

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

create

Keyword to create a meter.

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]

no adaptation-rule

Context

config>qos>access-ingress>meter

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command defines the method used by the system to derive the operational CIR and PIR rates when the meter is provisioned in hardware. For the **cir** and **pir** parameters, 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 CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default

adaptation-rule cir closest pir closest

Parameters

cir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced to adapt the CIR rate defined using the **meter meter-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used to derive the operational CIR rate for the meter. When

the **cir** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

Values **max** — Specifies that the operational CIR value is equal to or less than the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational CIR value is equal to or greater than the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

pir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced to adapt the PIR rate defined using the **meter meter-id rate** command. The **pir** parameter requires a qualifier that defines the constraint used to derive the operational PIR rate for the meter. When the **rate** command is not specified, the default constraint applies. The **max** (maximum), **min** (minimum), and **closest** qualifiers are mutually exclusive. See [Table 33: Supported hardware rates and burst step sizes for CIR and PIR values on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-T](#) for information about supported hardware step-size rates.

Default closest

Values **max** — Specifies that the operational PIR value is less than or equal to the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

min — Specifies that the operational PIR value is equal to or greater than the specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. The hardware step size varies based on the rate and the platform.

cbs

Syntax

cbs *size* [**kbits** | **bytes** | **kbytes**]

no cbs

Context

config>qos>access-ingress>meter

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command provides a mechanism to override the default committed burst size (CBS) for the meter. The *size* parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the CBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.

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

Default

cbs 32 kbits

Parameters

size

Specifies the size as an integer expression of the number of kilobits or kilobytes or bytes reserved for the meter. For example, if a value of 100 kbits is required, enter the value 100. The bucket size is rounded off to the next highest 4096 bytes boundary. The value can be specified in kilobits, kilobytes, or bytes.

Values	kbits — 4 to 2146959, default
	bytes — 512 to 274810752, default
	kbytes — 1 to 268369, default

color-mode

Syntax

color-mode {**color-aware** | **color-blind**}

no color-mode

Context

config>qos>access-ingress>meter

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command specifies whether the meter will operate in the **color-aware** or **color-blind** mode.

In **color-blind** mode, the profile and color assigned to the packet on ingress is ignored. The CIR and PIR rate configured for the meter is used to determine the final color and profile for the packet. If the packet is within the CIR, then the final profile and color assigned to the packet is in-profile and green. If the packet exceeds the CIR and is within the PIR, then the final profile and color assigned to the packet is out-of-profile and yellow. Packets that exceed the PIR rate are dropped.

In **color-aware** mode, the meter uses the profile assigned to the packet on ingress. The ingress profile can be assigned to the packet using either dot1p, DEI configured with dot1p, or IP DSCP values from the packet header.

In **color-aware** mode, the following behavior is expected.

- If the packet is pre-colored as an "in-profile" packet (also called "green" packets), depending on the burst size of the packet, the meter can mark the packet in-profile or out-profile.
- If the packet is pre-colored as an "out-profile" packet (which are also called "yellow" packets), even if the packet burst is less than the current available CBS, it would not be marked as in-profile and remain as out-profile.
- If the packet burst is higher than the MBS, it is marked as "red" and dropped by the meter at ingress.

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

Default

color-mode color-aware

Parameters

color-aware

Specifies that the meter will operate in color-aware mode.

color-blind

Specifies that the meter will operate in color-blind mode.

mbs

Syntax

mbs *size* [kbits | bytes | kbytes]

no mbs

Context

config>qos>access-ingress>meter

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command provides a mechanism to override the default MBS for the meter. The Maximum Burst Size (MBS) value specifies the maximum burst size that can be transmitted by the source while still complying

with the CIR. If the transmitted burst is lower than the configured MBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.

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

Default

512 kbits

Parameters

size

Specifies the size as integer expression of the number of kilobits reserved for the meter. For example, if a value of 100 KBits is required, enter the value 100. The bucket size is rounded off to the next highest 4096 bytes boundary.

Values **kbits** — 4 to 2146959, default
 bytes — 512 to 274810752, default
 kbytes — 1 to 268369, default

mode

Syntax

mode {**trtcm1** | **trtcm2** | **srtcm**}

no mode

Context

config>qos>access-ingress>meter

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command defines the mode of the meter. The mode can be configured as **trtcm1**, **trtcm2**, or **srtcm**. The **mode** command can be run at any time.

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

Default

mode trtcm1

Parameters

trtcm1

Keyword to implement the policing algorithm defined in RFC2698. Meters a packet stream and marks its packets either green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked either yellow or green depending on whether it exceeds

or does not exceed the CIR. The trTCM1 is useful, for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate.

trtcm2

Keyword to implement the policing algorithm defined in RFC4115. Meters a packet stream and marks its packets either green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked either yellow or green depending on whether it exceeds or does not exceed the CIR. The trtcm2 is useful, for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate.

srtcm

Keyword to implement the policing algorithm defined in RFC2697. Meters a packet stream and marks its packets either green, yellow, or red. Marking is based on a CIR and two associated burst sizes, a CBS and an MBS. A packet is marked green if it does not exceed the CBS, yellow if it exceeds the CBS but not the CIR, and red otherwise. The srTCM is useful, for example, for ingress policing of a service, where only the length, not the peak rate, of the burst determines service eligibility.

rate

Syntax

rate *cir* *cir-rate-in-kbps* [**pir** *pir-rate-in-kbps*]

no rate

Context

config>qos>access-ingress>meter

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command defines the administrative PIR and CIR parameters for the meter.

The **rate** command can be run at any time, altering the PIR and CIR rates for all meters created through the association of the network QoS policy with the meter ID.

The **no** form of this command reverts all meter instances created using this meter ID to the default PIR (max) and CIR parameters (0).



Note:

The value of rates are represented as 1000 kilobits per second and bursts are represented as 1024 kilobits.

Default

rate 0 pir max

Parameters

cir cir-rate-in-kbps

Specifies the administrative CIR rate, in kilobits, for the meter. The **cir** parameter overrides the default administrative CIR used by the meter. 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.

The actual CIR rate depends on the meter's **adaptation-rule** parameters and the hardware.

Values 0 to 40000000, max (7210 SAS-Sx/S 1/10GE)
0 to 64000000, max (7210 SAS-Mxp)
0 to 400000000, max (7210 SAS-Sx 10/100GE)
0 to 100000000, max (7210 SAS-R6 and 7210 SAS-R12)

Default 0

pir pir-rate-in-kbps

Specifies the administrative PIR rate, in kilobits, for the meter. When this parameter is configured, a valid PIR setting must be explicitly defined. If the **rate** command has not been run, the default PIR of **max** is assumed. If the **rate** command is run, a PIR setting is optional. The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir pir-rate-in-kbps** value.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate depends on the meter's **adaptation-rule** parameters and the hardware.



Note:

If the meter mode is configured as **trtcm2**, the system configures the policer EIR rate based on the PIR rate configured by the user.

Values 0 to 40000000, max (7210 SAS-Sx/S 1/10GE)
0 to 64000000, max (7210 SAS-Mxp)
0 to 400000000, max (7210 SAS-Sx 10/100GE)
0 to 100000000, max (7210 SAS-R6 and 7210 SAS-R12)

Default max

num-qos-classifiers

Syntax

num-qos-classifiers *num-resources*

no num-qos-classifiers

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command configures the number of CAM meter or policer resources that are allocated to rate-limit the forwarding class (FC) and classification entries to map the FC and traffic type to the configured meter. The maximum number of meters available for use by the FC defined under the policy is equal to half the value specified for the *num-resources* parameter. These meters are available for policing unicast or multipoint traffic, and for use by more than one FC.



Note:

- A user cannot modify the *num-resources* parameter when it is in use (that is, when the policy is applied to the port). To modify the meter or policer resources available for the policy, remove the association of the policy with the port, change the parameter to the desired value, and associate the policy with the port again.
- See section [Resource allocation for non-shared access ingress QoS policies](#) for information about resource allocation for access ingress QoS policies.

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

Default

num-qos-classifiers 4

Parameters

num-resources

Specifies the number of resources planned for use by the configured access ingress policy, expressed as a multiple of 2.

Values 4 to 32
 4 to 252 (only on the 7210 SAS-Mxp)

table-classification-criteria

Syntax

table-classification-criteria *table-classification-criteria*

Context

config>qos>access-ingress

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command allows the user to choose the classification policies to classify traffic to an FC when table-based classification is in use.

The following options can be used to configure the classification policy.

- If the **none** option is configured
 - use **default-fc** *fc-name* **profile** {**in** | **out**} (from the access ingress policy)
- If the **use-dscp** option is configured
 - use the DSCP classification policy for IP packets
 - use **default-fc** *fc-name* **profile** {**in** | **out**} (from the access ingress policy) for non-IP packets
- If the **use-dot1p** option is configured
 - use the dot1p classification policy for all tagged packets (IP and non-IP)
 - use **default-fc** *fc-name* **profile** {**in** | **out**} (from the access ingress policy) for untagged packets
- If the **both-dscp-dot1p** option is configured
 - use the DSCP classification policy for IP packets
 - use the dot1p classification policy for non-IP tagged packets
 - use **default-fc** *fc-name* **profile** {**in** | **out**} (from the access ingress policy) for non-IP untagged traffic

Default

table-classification-criteria both-dscp-dot1p

Parameters

table-classification-criteria

Specifies the table classification criteria to use.

Values use-dscp, use-dot1p, both-dscp-dot1p, none

20.2.2.3 Show commands

access-ingress

Syntax

access-ingress [*policy-id*] **association**

access-ingress [*policy-id*] **[detail]**

Context

show>qos

Platforms

7210 SAS-Mxp

Description

This command displays access ingress QoS policy information.

Parameters


- policy-id*
Displays information for the specified existing access ingress QoS policy.
- association**
Displays associations related to the access ingress QoS policy.
- detail**
Displays detailed access ingress QoS policy information.

Output

The following outputs are examples of access ingress QoS policy information, and the associated table describes the output fields:

- [Sample output \(shared access ingress QoS policy 61\)](#), [Table 81: Output fields: access ingress QoS policy \(shared and non-shared\)](#)
- [Sample output \(shared access ingress QoS policy 65536\)](#), [Table 81: Output fields: access ingress QoS policy \(shared and non-shared\)](#)
- [Sample output \(non-shared access ingress QoS policy 1\)](#), [Table 81: Output fields: access ingress QoS policy \(shared and non-shared\)](#)

Sample output (shared access ingress QoS policy 61)

 **Note:** For conciseness, some classificaiton entries have been removed.

```
*A:Mxp# show qos access-ingress 61

=====
QoS Access Ingress
=====
-----
Access Ingress Policy (61)
-----
Policy-id           : 61           Counter Mode       : forward-drop-count
Default FC          : ef           Profile             : In
DSCP Class Policy Id: 11           DOT1P Class Policy Id: 11
Share resources      : True         Classification crite*: cam
Table Classification criteria : both-dscp-dot1p
Description          : (Not Specified)
Criteria Type        : IP
Classifiers Allowed  : 132          Meters Allowed     : 16
-----
```

Table-based Resource Requirement									
Classifiers Reqrd		: 9		Meters Reqrd		: 8			

Cam-based Resource Requirement									
Classifiers Reqrd		: 31		Meters Reqrd		: 2			
=====									
* indicates that the corresponding row element may have been truncated.									
*A:Mxp# show qos access-ingress 61 detail									
=====									
QoS Access Ingress									
=====									
Access Ingress Policy (61)									

Policy-id		: 61		Counter Mode		: forward-drop-count			
Default FC		: ef		Profile		: In			
DSCP Class Policy Id:		11		DOT1P Class Policy Id:		11			
Share resources		: True		Classification crite*:		cam			
Table Classification criteria : both-dscp-dot1p									
Description : (Not Specified)									
Criteria Type		: IP							
Classifiers Allowed		: 132		Meters Allowed		: 16			

Table-based Resource Requirement									
Classifiers Reqrd		: 9		Meters Reqrd		: 8			

Cam-based Resource Requirement									
Classifiers Reqrd		: 31		Meters Reqrd		: 2			

Meter	Mode	CIR Admin	CIR Rule	PIR Admin	PIR Rule	CBS Admin	MBS Admin		
	Color Mode	CIR Oper		PIR Oper		CBS Oper	MBS Oper		

1	SrTcm	80000	closest	-	closest	5120 kbytes	5120 kbytes		
	color-aware	80000		-		5000 kbytes	5000 kbytes		
2	SrTcm	81000	closest	-	closest	10024 kbyt*	10024 kbyt*		
	color-aware	81024		-		9792 kbytes	9792 kbytes		
3	SrTcm	82000	closest	-	closest	4096 kbits	4096 kbits		
	color-aware	82000		-		4000 kbits	4000 kbits		
4	TrTcm1	50000	closest	83000	closest	100000 byt*	200000 byt*		
	color-aware	50000		83000		97792 bytes	195584 byt*		
5	TrTcm1	45000	closest	84000	closest	250000 kbi*	250000 kbi*		
	color-blind	45056		83968		244160 kbi*	244160 kbi*		
6	TrTcm1	55000	closest	85000	closest	90000 kbits	10000 kbits		
	color-blind	54976		85000		87904 kbits	9768 kbits		
7	TrTcm1	56000	closest	86000	closest	80000 bytes	90000 bytes		
	color-aware	56000		86000		78336 bytes	88064 bytes		
8	TrTcm2	37000	closest	50000	closest	8092 kbytes	9012 kbytes		
	color-aware	36992		49984		7904 kbytes	8804 kbytes		
9	TrTcm2	37000	closest	51000	closest	10024 kbits	10082 kbits		
	color-aware	37000		51000		9792 kbits	9848 kbits		
10	TrTcm2	37000	closest	52000	closest	def kbits	def kbits		
	color-aware	37000		52000		32 kbits	1000 kbits		
11	TrTcm1	38000	closest	90000	closest	20048 kbits	40048 kbits		

12	color-aware	38000		90016		19584 kbits	39120 kbits
	TrTcm2	45000	closest	46000	closest	2100000 by*	2100000 by*
	color-aware	45000		46000		2051072 by*	2051072 by*
13	TrTcm1	45000	closest	92000	closest	9000 kbits	9000 kbits
	color-aware	45000		92000		8792 kbits	8792 kbits
14	TrTcm1	50000	closest	93000	closest	10000 kbits	10000 kbits
	color-aware	50000		93000		9768 kbits	9768 kbits
15	TrTcm1	51000	closest	94000	closest	def kbits	def kbits
	color-aware	51000		94000		32 kbits	1000 kbits
16	TrTcm2	52000	closest	43000	closest	100034 byt*	100044 kby*
	color-aware	52000		43008		97792 bytes	97728 kbyt*

FC	UCastM	MCastM					

be	2	def					
l2	7	def					
af	def	def					
l1	6	def					
h2	5	def					
ef	3	def					
h1	4	def					
nc	8	def					

Port Attachments							

Port-id : 1/1/12							
Port-id : 1/1/13							
Port-id : 1/1/14							
Port-id : lag-1							

Match Criteria							

IP Match Criteria							

Entry	: 1						
Description	: (Not Specified)						
Source IP	: Undefined						
Dest. IP	: Undefined						
Source Port	: None			Dest. Port	: None		
Protocol	: none			DSCP	: None		
FC	: af			Ip Precedence	: None		
=====							
* indicates that the corresponding row element may have been truncated.							

Sample output (shared access ingress QoS policy 65536)

```
*A:Mxp# show qos access-ingress 65536

=====
QoS Access Ingress
=====
Access Ingress Policy (65536)
-----
Policy-id          : 65536          Counter Mode       : in-out-profile-count
Default FC         : be             Profile            : Out
DSCP Class Policy Id: 1             DOT1P Class Policy Id: 1
Share resources    : True           Classification crite*: table
```

```

Table Classification criteria : both-dscp-dot1p
Description      : Default ACCESS ingress shared QoS policy.
Criteria Type    : None
Classifiers Allowed : 2          Meters Allowed      : 1

-----
Table-based Resource Requirement
-----
Classifiers Reqrd : 2          Meters Reqrd        : 1

-----
Cam-based Resource Requirement
-----
Classifiers Reqrd : 1          Meters Reqrd        : 1

=====
* indicates that the corresponding row element may have been truncated.

*A:Mxp# show qos access-ingress 65536 detail

=====
QoS Access Ingress
=====
-----
Access Ingress Policy (65536)
-----
Policy-id      : 65536          Counter Mode    : in-out-profile-count
Default FC     : be            Profile          : Out
DSCP Class Policy Id: 1        DOT1P Class Policy Id: 1
Share resources : True         Classification crite*: table
Table Classification criteria : both-dscp-dot1p
Description     : Default ACCESS ingress shared QoS policy.
Criteria Type   : None
Classifiers Allowed : 2          Meters Allowed      : 1

-----
Table-based Resource Requirement
-----
Classifiers Reqrd : 2          Meters Reqrd        : 1

-----
Cam-based Resource Requirement
-----
Classifiers Reqrd : 1          Meters Reqrd        : 1

-----
Meter Mode      CIR Admin CIR Rule PIR Admin PIR Rule CBS Admin MBS Admin
Color Mode     CIR Oper          PIR Oper          CBS Oper  MBS Oper
-----
1      TrTcm1      0          closest max      closest def kbits  def kbits
      color-aware 0          64000000 32 kbits  1024 kbits

-----
FC          UCastM      MCastM
-----
No FC-Map Entries Found.

-----
Port Attachments
-----
Port-id : 1/1/2

```

```
Port-id : 1/1/10
Port-id : 1/1/11
Port-id : lag-20
-----
Match Criteria
-----
No Matching Criteria.

=====
* indicates that the corresponding row element may have been truncated.
```

Sample output (non-shared access ingress QoS policy 1)

```
*A:Mxp# show qos access-ingress 1

=====
QoS Access Ingress
=====
-----
Access Ingress Policy (1)
-----
Policy-id           : 1           Counter Mode       : in-out-profile-count
Default FC          : be          Profile             : Out
DSCP Class Policy Id: 1           DOT1P Class Policy Id: 1
Share resources      : False       Classification crite*: table
Table Classification criteria : both-dscp-dot1p
Description          : Default ACCESS ingress QoS policy.
Criteria Type        : None
Classifiers Allowed  : 4           Meters Allowed     : 2
-----
Access Ingress Resource Requirement
-----
Filters Reqr        : 4           Meters Reqr        : 2
-----
* indicates that the corresponding row element may have been truncated.
```

Table 81: Output fields: access ingress QoS policy (shared and non-shared)

Label	Description
Policy-Id	Displays the ID of the policy
Counter Mode	in-out-profile-count — Displays the in-profile and out-profile packet count per meter forward-drop-count — Displays the forwarded and dropped packets count per meter
Default FC	Displays the default FC for the policy
Profile	Displays the profile configured
DSCP Class Policy Id	Displays the DSCP classification policy ID
DOT1P Class Policy Id	Displays the dot1p classification policy ID

Label	Description
Share Resources	true — Policy is configured to share resources false — Policy is not configured to share resources
Classification criteria	cam — Policy is configured with CAM-based classification criteria table — Policy is configured with table-based classification criteria
Table classification criteria	Displays the table classification criteria used for the policy
Description	Description of the policy
Criteria Type	Displays the type of cam-criteria configured (for example, IP criteria).
Classifiers Allowed	Displays the number of classifiers allowed for the policy
Meters Allowed	Displays the number of meters allowed for the policy
Meter	Displays the meter ID
Mode	Displays the configured mode of the meter (trTcm1 or srTcm)
Color Mode	Displays the color mode of the meter (color-blind or color-aware)
CIR Admin	Displays the administrative Committed Information Rate (CIR) parameters for the meters
CIR Oper	Displays the operational Committed Information Rate (CIR) parameters for the meters
CIR Rule	min — Operational CIR for the meters is equal to or greater than the administrative rate specified using the rate command max — Operational CIR for the meter is equal to or less than the administrative rate specified using the rate command closest — Operational PIR for the meters are the rate closest to the rate specified using the rate command, without exceeding the operational PIR
PIR Admin	Displays the administrative Peak Information Rate (PIR) parameters for the meters
PIR Oper	Displays the operational PIR parameters for the meter
PIR Rule	min — Operational PIR for the meter is equal to or greater than the administrative rate specified using the rate command max — Operational PIR for the meters is equal to or less than the administrative rate specified using the rate command

Label	Description
	closest — Operational PIR for the meters is the rate closest to the rate specified using the rate command
CBS Admin	def — Displays the default CBS value for the meters
CBS Oper	value — Displays the value to override the default CBS for the meters
MBS Admin	def — Displays the default MBS value
MBS Oper	value — Displays the value to override the default MBS for the meter
FC	Displays the forwarding class
UCastM	Displays the default unicast forwarding type meter mapping
MCastM	Displays the overrides for the default multicast forwarding type meter mapping
Entry	Displays the entry number for the IP match criteria
Source IP	Displays the source IP address for the IP match criteria
Dest. IP	Displays the destination IP address for the IP match criteria
Source Port	Displays the source port for the IP match criteria
Dest. Port	Displays the destination port for the IP match criteria
Protocol	Displays the protocol for the IP match criteria
DSCP	Displays the DSCP for the IP match criteria
FC	Displays the FC for the IP match criteria
IP Precedence	Displays the IP precedence for the IP match criteria

21 Standards and protocol support



Note:

The information provided in this chapter is subject to change without notice and may not apply to all platforms.

Nokia assumes no responsibility for inaccuracies.

The following conventions are used in this section:

- T(A,N) indicates 7210 SAS-T in both Access-uplink mode and Network mode. Similarly, T(N) indicates 7210 SAS-T in network mode only.
- K5 indicates 7210 SAS-K 2F2T1C.
- K12 indicates 7210 SAS-K 2F4T6C.
- K30 indicates 7210 SAS-K 3SFP+ 8C.
- Sx/S-1/10GE indicates all variants of 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms.
- Sx-10/100GE indicates only the variants of 7210 SAS-Sx 10/100GE. Sx-10/100GE QSFP28 indicates the 7210 SAS-Sx 10/100GE 64 SFP+ 4QSFP28 variant.
- Sx/S-1/10GE-VC indicates 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms in standalone VC mode.
- R6 indicates 7210 SAS-R6.
- R12 indicates 7210 SAS-R12.
- D indicates 7210 SAS-D and 7210 SAS-D ETR. If a line item applies only to 7210 SAS-D ETR, then it is indicated as D-ETR.
- Dxp indicates 7210 SAS-Dxp-12p (2SFP+ 4F6T) 7210 SAS-Dxp-12p ETR (2SFP+ 4F6T), 7210 SAS-Dxp 16p (2SFP+ 4F10T), and 7210 SAS-Dxp-24p (2SFP+ 6F16T). If a line item applies only to a particular variant, the variant name will be called out explicitly against that item.
- This standards list is not applicable to platforms in the satellite mode of operation, as most of the features are supported on 7x50 SR platforms. For this reason, the host platforms standards compliance must be consulted for the satellite mode of operation.

21.1 BGP

draft-ietf-idr-add-paths-04, Advertisement of Multiple Paths in BGP is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-sidr-origin-validation-signaling-04, BGP Prefix Origin Validation State Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1772, Application of the Border Gateway Protocol in the Internet is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1997, BGP Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2385, Protection of BGP Sessions via the TCP MD5 Signature Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2439, BGP Route Flap Damping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2858, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2918, Route Refresh Capability for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3107, Carrying Label Information in BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3392, Capabilities Advertisement with BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4271, A Border Gateway Protocol 4 (BGP-4) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4360, BGP Extended Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4456, BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only R6 and R12 supports RR server functionality. Rest of the platforms support only client function.

RFC 4659, BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4684, Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4760, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4798, Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4893, BGP Support for Four-octet AS Number Space is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5004, Avoid BGP Best Path Transitions from One External to Another is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5291, Outbound Route Filtering Capability for BGP-4 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5668, 4-Octet AS Specific BGP Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6811, Prefix Origin Validation is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

21.2 Ethernet

IEEE 802.1AB, Station and Media Access Control Connectivity Discovery is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

TIA-1057, LLDP for Media endpoint devices is supported on Dxp, Sx/S-1/10GE, and Sx/S-1/10GE-VC

IEEE 802.1ad, Provider Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ag, Connectivity Fault Management is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ah, Provider Backbone Bridges is supported on T(N)

IEEE 802.1ax, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1D, MAC Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1p, Traffic Class Expediting is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1Q, Virtual LANs is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1s, Multiple Spanning Trees is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1w, Rapid Reconfiguration of Spanning Tree is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1X, Port Based Network Access Control is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ab, 1000BASE-T is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ac, VLAN Tag is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ad, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ae, 10 Gb/s Ethernet is supported on Dxp, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ah, Ethernet in the First Mile is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ba, 40 Gb/s and 100 Gb/s Ethernet is supported on R6, R12, and Sx-10/100GE

IEEE 802.3i, Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3u, Fast Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3z, Gigabit Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3af, Power Over Ethernet (PoE) is supported on Dxp, T-ETR, Mxp-ETR, and Sx/S-1/10GE



Note:

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3at, Power Over Ethernet (PoE+) is supported on Dxp, T-ETR, Mxp-ETR, and Sx/S-1/10GE



Note:

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3bt, Power Over Ethernet (PoE++/HPoE) is supported on Dxp



Note:

Only on Dxp-16p and Dxp-24p.

ITU-T G.8032, Ethernet Ring Protection Switching is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ITU-T Y.1731, OAM functions and mechanisms for Ethernet based networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

21.3 EVPN

draft-snr-bess-evpn-proxy-arp-nd-00, Proxy-ARP/ND function in EVPN networks is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:

Sx/S-1/10GE standalone mode only.

RFC 7432, BGP MPLS-Based Ethernet VPN is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:

Sx/S-1/10GE standalone mode only.

draft-ietf-bess-evpn-vpws-14, Virtual Private Wire Service support in Ethernet VPN is supported on Mxp

21.4 Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, Operational management of Loop Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

With Segment Routing.

RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

With Segment Routing.

RFC 7490, Remote Loop-Free Alternate (LFA) Fast Reroute (FRR) is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

With Segment Routing.

21.5 Internet Protocol (IP) — General

draft-grant-tacacs-02, The TACACS+ Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-rrp-unified-spec-02, Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6 is supported on Mxp

RFC 768, User Datagram Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 793, Transmission Control Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 854, TELNET Protocol Specifications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 951, Bootstrap Protocol (BOOTP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1034, Domain Names - Concepts and Facilities is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1035, Domain Names - Implementation and Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1350, The TFTP Protocol (revision 2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1534, Interoperation between DHCP and BOOTP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1542, Clarifications and Extensions for the Bootstrap Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2131, Dynamic Host Configuration Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2132, DHCP Options and BOOTP Vendor Extensions is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2347, TFTP Option Extension is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2348, TFTP Blocksize Option is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2349, TFTP Timeout Interval and Transfer Size Options is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2428, FTP Extensions for IPv6 and NATs is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2865, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2866, RADIUS Accounting is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3046, DHCP Relay Agent Information Option (Option 82) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3596, DNS Extensions to Support IP version 6 is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3768, Virtual Router Redundancy Protocol (VRRP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4251, The Secure Shell (SSH) Protocol Architecture is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4252, The Secure Shell (SSH) Authentication Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

All 7210 platforms support password and publickey based user authentication. 7210 SAS-D support only password based authentication.

RFC 4253, The Secure Shell (SSH) Transport Layer Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4254, The Secure Shell (SSH) Connection Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4632, Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5880, Bidirectional Forwarding Detection (BFD) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5881, Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 6528, Defending against Sequence Number Attacks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7130, Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

21.6 IP — Multicast

RFC 1112, Host Extensions for IP Multicasting is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IGMP v1, v2, v3 is supported.

RFC 2236, Internet Group Management Protocol, Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3306, Unicast-Prefix-based IPv6 Multicast Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3376, Internet Group Management Protocol, Version 3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3446, Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

MSDP supported only on Sx/S-1/10GE standalone.

RFC 3618, Multicast Source Discovery Protocol (MSDP) is supported on Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4604, Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

MLD not supported.

RFC 4607, Source-Specific Multicast for IP is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4608, Source-Specific Protocol Independent Multicast in 232/8 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5059, Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5384, The Protocol Independent Multicast (PIM) Join Attribute Format is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6513, Multicast in MPLS/BGP IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6514, BGP Encodings and Procedures for Multicast in MPLS/IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6515, IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6625, Wildcards in Multicast VPN Auto-Discover Routes is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6826, Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7246, Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7385, IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

21.7 IP — Version 4

RFC 791, Internet Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 792, Internet Control Message Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 826, An Ethernet Address Resolution Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1519, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1812, Requirements for IPv4 Routers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1981, Path MTU Discovery for IP version 6 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2401, Security Architecture for Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Supported only for OSPFv3 authentication. Not supported for services.

RFC 2460, Internet Protocol, Version 6 (IPv6) Specification is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

21.8 IP — Version 6

RFC 2464, Transmission of IPv6 Packets over Ethernet Networks is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 3021, Using 31-Bit Prefixes on IPv4 Point-to-Point Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3122, Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification is supported on K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3587, IPv6 Global Unicast Address Format is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4007, IPv6 Scoped Address Architecture is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4193, Unique Local IPv6 Unicast Addresses is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4291, Internet Protocol Version 6 (IPv6) Addressing Architecture is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4861, Neighbor Discovery for IP version 6 (IPv6) is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5095, Deprecation of Type 0 Routing Headers in IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5952, A Recommendation for IPv6 Address Text Representation is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 6106, IPv6 Router Advertisement Options for DNS Configuration is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6164, Using 127-Bit IPv6 Prefixes on Inter-Router Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

21.9 IPsec

RFC 2401, Security Architecture for the Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only for use with OSPFv3 authentication. Not supported for services.

RFC 2406, IP Encapsulating Security Payload (ESP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only for use with OSPFv3 authentication. Not supported for services.

21.10 IS-IS

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-kaplan-isis-ext-eth-02, Extended Ethernet Frame Size Support is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

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) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3359, Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3719, Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3787, Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4971, Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5120, M-ISIS: Multi Topology (MT) Routing in IS-IS is supported on K12, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5130, A Policy Control Mechanism in IS-IS Using Administrative Tags is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5301, Dynamic Hostname Exchange Mechanism for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5302, Domain-wide Prefix Distribution with Two-Level IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5304, IS-IS Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5305, IS-IS Extensions for Traffic Engineering TE is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5306, Restart Signaling for IS-IS (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5308, Routing IPv6 with IS-IS is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5310, IS-IS Generic Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6232, Purge Originator Identification TLV for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6233, IS-IS Registry Extension for Purges is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

K12, K30 support only a single instance and can operate in multi-instance deployment as it supports the processing of TLVs for multi-instance support.

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

21.11 Management

draft-ietf-snmpv3-update-mib-05, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-idr-bgp4-mib-05, Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-wg-mib-06, Management Information Base for Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mboned-msdp-mib-01, Multicast Source Discovery protocol MIB is supported on Sx/S-1/10GE



Note:

Only in standalone mode.

draft-ietf-mpls-ldp-mib-07, Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mpls-lsr-mib-06, Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mpls-te-mib-04, Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-mib-update-08, OSPF Version 2 Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaaddressfamilynumbers-mib, IANA-ADDRESS-FAMILY-NUMBERS-MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaifttype-mib, IANAifType-MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaiprouteprotocol-mib, IANA-RTPROTO-MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8021-CFM-MIB, IEEE P802.1ag(TM) CFM MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8021-PAE-MIB, IEEE 802.1X MIB is supported on D, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8023-LAG-MIB, IEEE 802.3ad MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

LLDP-MIB, IEEE P802.1AB(TM) LLDP MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1157, A Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1215, A Convention for Defining Traps for use with the SNMP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1724, RIP Version 2 MIB Extension is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2021, Remote Network Monitoring Management Information Base Version 2 using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2138, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2206, RSVP Management Information Base using SMIv2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2213, Integrated Services Management Information Base using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2571, An Architecture for Describing SNMP Management Frameworks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2573, SNMP Applications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2575, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2578, Structure of Management Information Version 2 (SMIv2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2579, Textual Conventions for SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2787, Definitions of Managed Objects for the Virtual Router Redundancy Protocol is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2819, Remote Network Monitoring Management Information Base is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2856, Textual Conventions for Additional High Capacity Data Types is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2863, The Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2864, The Inverted Stack Table Extension to the Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2933, Internet Group Management Protocol MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3014, Notification Log MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3164, The BSD syslog Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3165, Definitions of Managed Objects for the Delegation of Management Scripts is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3231, Definitions of Managed Objects for Scheduling Management Operations is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3273, Remote Network Monitoring Management Information Base for High Capacity Networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3414, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3416, Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3417, Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3419, Textual Conventions for Transport Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3584, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3593, Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3635, Definitions of Managed Objects for the Ethernet-like Interface Types is supported on D, E, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3826, The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3877, Alarm Management Information Base (MIB) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4001, Textual Conventions for Internet Network Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4022, Management Information Base for the Transmission Control Protocol (TCP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4113, Management Information Base for the User Datagram Protocol (UDP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4220, Traffic Engineering Link Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4292, IP Forwarding Table MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4293, Management Information Base for the Internet Protocol (IP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5101, Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information is supported on Mxp, Sx/S-1/10GE, and R6

RFC 6241, Network Configuration Protocol (NETCONF) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

RFC 6242, Using the NETCONF Protocol over Secure Shell (SSH) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

21.12 MPLS — General

RFC 3031, Multiprotocol Label Switching Architecture is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3032, MPLS Label Stack Encoding is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3443, Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4182, Removing a Restriction on the use of MPLS Explicit NULL is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5332, MPLS Multicast Encapsulations is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

21.13 MPLS — GMPLS

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

21.14 MPLS — LDP

draft-ietf-mpls-ldp-ipv6-15, Updates to LDP for IPv6 is supported on Mxp

draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-v2-00, LDP Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-hello-reduce-04, Targeted LDP Hello Reduction is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3037, LDP Applicability is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5036, LDP Specification is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5443, LDP IGP Synchronization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5561, LDP Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:

P2MP LSPs only.

21.15 MPLS — MPLS-TP

RFC 5586, MPLS Generic Associated Channel is supported on T(N), R6, and R12

RFC 5921, A Framework for MPLS in Transport Networks is supported on T(N), R6, and R12

RFC 5960, MPLS Transport Profile Data Plane Architecture is supported on T(N), R6, and R12

RFC 6370, MPLS Transport Profile (MPLS-TP) Identifiers is supported on T(N), R6, and R12

RFC 6378, MPLS Transport Profile (MPLS-TP) Linear Protection is supported on T(N), R6, and R12

RFC 6426, MPLS On-Demand Connectivity and Route Tracing is supported on T(N), R6, and R12

RFC 6428, Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile is supported on T(N), R6, and R12

RFC 6478, Pseudowire Status for Static Pseudowires is supported on T(N), R6, and R12

RFC 7213, MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing is supported on T(N), R6, and R12

21.16 MPLS — OAM

RFC 6424, Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6425, Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping is supported on T(N), Mxp, R6, and R12

21.17 MPLS — RSVP-TE

RFC 2702, Requirements for Traffic Engineering over MPLS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2747, RSVP Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2961, RSVP Refresh Overhead Reduction Extensions is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3097, RSVP Cryptographic Authentication -- Updated Message Type Value is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3477, Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4561, Definition of a Record Route Object (RRO) Node-Id Sub-Object is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4875, Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs) is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4950, ICMP Extensions for Multiprotocol Label Switching is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5817, Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

21.18 OSPF

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1765, OSPF Database Overflow is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2328, OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3101, The OSPF Not-So-Stubby Area (NSSA) Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3509, Alternative Implementations of OSPF Area Border Routers is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3623, Graceful OSPF Restart Graceful OSPF Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3630, Traffic Engineering (TE) Extensions to OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4222, Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4552, Authentication/Confidentiality for OSPFv3 is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4576, Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4970, Extensions to OSPF for Advertising Optional Router Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5185, OSPF Multi-Area Adjacency is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5187, OSPFv3 Graceful Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5243, OSPF Database Exchange Summary List Optimization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5250, The OSPF Opaque LSA Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5340, OSPF for IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5838, Support of Address Families in OSPFv3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6987, OSPF Stub Router Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

21.19 Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, OAM Procedures for VPWS Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3916, Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3985, Pseudo Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4385, Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4446, IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4447, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5659, An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6073, Segmented Pseudowire is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6310, Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6391, Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network is supported on K12, K30, Mxp, R6, and R12

RFC 6718, Pseudowire Redundancy is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6870, Pseudowire Preferential Forwarding Status bit is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7023, MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7267, Dynamic Placement of Multi-Segment Pseudowires is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

21.20 Quality of Service

RFC 2430, A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2598, An Expedited Forwarding PHB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3140, Per Hop Behavior Identification Codes is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3260, New Terminology and Clarifications for Diffserv is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

21.21 RIP

RFC 1058, Routing Information Protocol is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2082, RIP-2 MD5 Authentication is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2453, RIP Version 2 is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

21.22 Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, IEEE default profile is supported only includes the Dxp-12p ETR, Dxp-16p, Dxp-24p. Sx-10/100GE does not support IEEE default profile.

ITU-T G.781, Synchronization layer functions, issued 09/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8261, Timing and synchronization aspects in packet networks, issued 04/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8262, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8264, Distribution of timing information through packet networks, issued 10/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8265.1, Precision time protocol telecom profile for frequency synchronization, issued 10/2010 is supported on D-ETR, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, R6, and R12

ITU-T G.8275.1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014 is supported on Dxp, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12



Note:

For 7210 SAS-Sx 10/100GE, the support only includes the Sx 10/100GE QSFP28 variant. For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

IEC/IEEE 61850-9-3-2016, Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation is supported on Dxp-16p and Dxp-24p

IEEE C37.238-2017 - IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications is supported on Dxp-16p and Dxp-24p

RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

21.23 VPLS

RFC 4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

On 7210 platforms, only BGP-AD is supported with TLDP signalling for PW. No BGP signalling is supported for PW establishment.

RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5501, Requirements for Multicast Support in Virtual Private LAN Services is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)