



7210 Service Access System

Release 26.9.R1

7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Quality of Service Guide

3HE 27840 AAAB TQZZA 01
Edition: 01
September 2026

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.

© 2026 Nokia.

Table of contents

List of tables.....	12
List of figures.....	15
1 Getting started.....	16
1.1 About this guide.....	16
1.1.1 Document structure and content.....	16
1.2 7210 SAS modes of operation.....	17
1.3 7210 SAS port modes.....	19
1.4 7210 SAS QoS configuration process.....	21
1.5 Conventions.....	22
1.5.1 Precautionary and information messages.....	22
1.5.2 Options or substeps in procedures and sequential workflows.....	23
2 QoS policies.....	24
2.1 QoS policies overview.....	24
2.1.1 Overview of QoS policies on 7210 SAS-K 2F1C2T.....	25
2.1.2 Overview of QoS policies on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	27
2.1.2.1 Overview of network QoS policy applied to access-uplink ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	30
2.1.2.2 Overview of network QoS policy applied to network and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	31
2.1.3 Summary of major functions of QoS policies.....	32
2.2 Network and service QoS policies.....	35
2.2.1 Network QoS policies on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C network ports and hybrid ports.....	35
2.2.1.1 Ingress classification support for network ports and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	36
2.2.1.2 Egress marking support for network ports and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	37
2.2.2 Network QoS policies on access-uplink ports.....	39
2.2.2.1 Ingress classification support for access-uplink ports.....	39
2.2.3 Network queue policies.....	41
2.2.4 Service ingress QoS policies.....	44
2.2.4.1 Default service ingress policy.....	45

2.2.5	Service ingress classification.....	46
2.2.5.1	Service ingress classification.....	46
2.2.5.2	Service ingress classification — IP and MAC packet fields.....	46
2.2.6	Service egress QoS policies.....	51
2.2.6.1	Default service egress policy.....	52
2.3	Meters/policers.....	52
2.3.1	Meter/policers and policer parameters.....	52
2.3.1.1	Meter ID.....	53
2.3.1.2	Committed information rate.....	53
2.3.1.3	Peak information rate.....	53
2.3.1.4	Adaptation rule for meters.....	53
2.3.1.5	Committed burst size.....	54
2.3.1.6	Maximum burst size.....	54
2.3.1.7	Meter counters.....	55
2.3.1.8	Meter modes.....	55
2.3.1.9	Meter platform support.....	55
2.3.2	Ingress profile assignment.....	55
2.3.3	QoS override.....	56
2.3.3.1	Configuring meter override parameters.....	57
2.4	FCs.....	57
2.4.1	Forwarding-class to queue ID mapping.....	58
2.4.2	FC to queue ID mapping.....	58
2.5	Schedulers.....	58
2.6	CPU queues.....	59
2.7	Egress port rate limiting.....	59
2.8	Queue management.....	59
2.8.1	Queues and queue parameters.....	59
2.8.1.1	Queue ID.....	60
2.8.1.2	Committed information rate.....	60
2.8.1.3	Peak information rate.....	61
2.8.1.4	Adaptation rule for queues.....	61
2.8.1.5	Committed burst size.....	62
2.8.1.6	Maximum burst size.....	62
2.8.1.7	Ingress profile assignment.....	62
2.8.1.8	Weight and priority.....	63
2.8.1.9	Queue counters.....	63

2.8.1.10	Queue platform support.....	63
2.8.2	Buffer pools.....	64
2.8.2.1	Ring and non-ring buffer pool.....	66
2.8.2.2	Configuration guidelines for CBS and MBS.....	68
2.8.3	Slope policies.....	69
2.8.4	RED slopes.....	69
2.8.4.1	Operation and configuration of RED slopes.....	69
2.8.4.2	Simplified overview of RED.....	69
2.8.4.3	Slope policy parameters.....	71
2.9	Preclassification on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	71
2.10	QoS policy entities.....	72
2.10.1	Summary of QoS policy support for hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	72
2.11	Configuration notes.....	73
3	DEI-based classification and marking.....	74
3.1	DEI-based classification and marking.....	74
3.2	DEI-based marking.....	74
4	Port-level egress rate limiting.....	75
4.1	Overview.....	75
4.2	Basic configurations.....	76
4.2.1	Modifying the port-level egress-rate command.....	76
4.2.2	Removing the port-level egress-rate command.....	76
4.3	Port level egress-rate command reference.....	77
4.3.1	Command hierarchies.....	77
4.3.1.1	Configuration commands.....	77
4.3.1.2	Show commands.....	77
4.3.2	Command descriptions.....	77
4.3.2.1	Configuration commands.....	77
4.3.2.2	Show commands.....	78
5	Frame-based accounting.....	86
5.1	Overview of frame-based accounting.....	86
5.2	Enabling and disabling frame-based accounting.....	86
5.3	Frame-based accounting command reference.....	87
5.3.1	Command hierarchies.....	87

5.3.1.1	Configuration commands.....	87
5.3.1.2	Show commands.....	87
5.3.2	Command descriptions.....	87
5.3.2.1	Configuration commands.....	88
5.3.2.2	Show commands.....	88
6	DSCP, dot1p, and MPLS EXP classification policies.....	94
6.1	Overview.....	94
6.1.1	DSCP classification policy.....	94
6.1.2	dot1p classification policy.....	94
6.1.3	MPLS EXP classification policy.....	94
6.2	Configuring classification policies.....	95
6.3	DSCP, dot1p, and MPLS EXP classification policy command reference.....	96
6.3.1	Command hierarchies.....	96
6.3.1.1	Configuration commands for 7210 SAS-K 2F1C2T.....	96
6.3.1.2	Configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	96
6.3.1.3	Show commands.....	97
6.3.1.4	Operational commands.....	97
6.3.2	Command description.....	97
6.3.2.1	Configuration commands.....	97
6.3.2.2	Show commands.....	105
7	Network QoS policies.....	111
7.1	Overview of network QoS policy on 7210 SAS-K 2F1C2T.....	111
7.1.1	Resource allocation for network QoS policy for 7210 SAS-K 2F1C2T.....	112
7.2	Basic configurations for 7210 SAS-K 2F1C2T.....	112
7.2.1	Create a network QoS policy for access-uplink ports on 7210 SAS-K 2F1C2T.....	112
7.2.1.1	Default network policy values for access-uplink ports on 7210 SAS-K 2F1C2T..	114
7.3	Overview of network QoS policy on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	117
7.3.1	Network QoS policy for access-uplink ports on 7210 SAS-K 2F6C4T and 7210 SAS- K 3SFP+ 8C.....	117
7.3.2	Network QoS policy for network ports and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	118
7.3.3	Resource allocation for network QoS policy for 7210 SAS-K 2F6C4T and 7210 SAS- K 3SFP+ 8C.....	119
7.4	Basic configurations for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	119

7.4.1	Create a network QoS policy on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C....	119
7.4.1.1	Create a network QoS policy for access-uplink ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	119
7.4.1.2	Create a network QoS policy for network ports and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	123
7.4.1.3	Default network policy values on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	126
7.5	DSCP marking CPU self-generated traffic.....	130
7.5.1	QoS for CPU self-generated traffic on network interfaces for the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	133
7.5.2	Default DSCP mapping.....	133
7.6	Service management tasks.....	134
7.6.1	Deleting QoS policies.....	134
7.6.2	Remove a policy from the QoS configuration.....	134
7.6.3	Copying and overwriting network policies.....	135
7.6.4	Editing QoS policies.....	135
7.7	Network QoS policy command reference.....	135
7.7.1	Command hierarchies.....	135
7.7.1.1	Configuration commands for 7210 SAS-K 2F1C2T.....	136
7.7.1.2	Configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	136
7.7.1.3	Self-generated traffic commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	137
7.7.1.4	Operational commands.....	137
7.7.1.5	Show commands.....	138
7.7.2	Command descriptions.....	138
7.7.2.1	Configuration commands.....	138
7.7.2.2	Show commands.....	159
8	Network queue QoS policies.....	171
8.1	Overview.....	171
8.2	Basic configurations.....	171
8.2.1	Creating a network queue QoS policy.....	171
8.2.2	Applying network queue policies.....	172
8.2.2.1	Applying network queue configuration to access-uplink ports.....	172
8.2.2.2	Applying network queue configuration to network ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	173
8.3	Default network queue policy values.....	173

8.3.1	Default network queue policy for 7210 SAS-K 2F1C2T.....	174
8.3.2	Default network queue policy for 7210 SAS-K 2F6C4T.....	175
8.3.3	Default network queue policy for 7210 SAS-K 3SFP+ 8C.....	177
8.4	Service management tasks.....	179
8.4.1	Deleting network queue QoS policies.....	179
8.4.2	Copying and overwriting network queue QoS policies.....	179
8.4.3	Editing network queue QoS policies.....	180
8.5	Network queue QoS policy command reference.....	181
8.5.1	Command hierarchies.....	181
8.5.1.1	Configuration commands.....	181
8.5.1.2	Operational commands.....	181
8.5.1.3	Show commands.....	181
8.5.2	Command descriptions.....	182
8.5.2.1	Configuration commands.....	182
8.5.2.2	Show commands.....	190
9	Service ingress QoS policies.....	194
9.1	Overview of service ingress policy.....	194
9.1.1	Configuration guidelines for SAP-ingress policy.....	194
9.1.1.1	Resource allocation for service ingress QoS classification policy.....	196
9.1.1.2	Resource allocation for SAP ingress meters.....	196
9.1.1.3	Default SAP-ingress policy.....	197
9.1.1.4	Use of index file by SAP QoS ingress policy.....	198
9.2	Basic configurations.....	199
9.2.1	Service ingress QoS policies.....	199
9.2.2	Service ingress QoS queues.....	200
9.2.3	Service ingress QoS meters.....	200
9.2.4	SAP ingress FC configuration.....	201
9.2.5	Service ingress dot1p and IP DSCP criteria.....	202
9.2.6	Service ingress IP match criteria.....	202
9.2.7	Service ingress MAC match criteria.....	203
9.2.8	Applying service ingress policies.....	204
9.2.8.1	Epipe service.....	204
9.2.8.2	VPLS.....	204
9.2.8.3	IES.....	205
9.3	Service management tasks.....	205

9.3.1	Deleting QoS policies.....	205
9.3.1.1	Removing a QoS policy from service SAPs.....	205
9.3.2	Copying and overwriting QoS policies.....	206
9.3.3	Removing a policy from the QoS configuration.....	207
9.3.4	Editing QoS policies.....	207
9.4	Service SAP QoS policy command reference.....	207
9.4.1	Command hierarchies.....	207
9.4.1.1	Service ingress QoS policy commands.....	207
9.4.1.2	Operational commands.....	209
9.4.1.3	Show commands.....	209
9.4.2	Command descriptions.....	209
9.4.2.1	Configuration commands.....	210
9.4.2.2	Operational commands.....	211
9.4.2.3	Service ingress QoS policy commands.....	217
9.4.2.4	Service ingress QoS policy entry commands.....	227
9.4.2.5	Service ingress MAC QoS policy match commands.....	232
9.4.2.6	SAP ingress queue and meter QoS policy commands.....	238
9.4.2.7	Show commands.....	249
10	Service egress QoS policies.....	257
10.1	Overview.....	257
10.1.1	Egress SAP FC and profile overrides.....	258
10.1.2	Configuration guidelines for access SAP egress policies.....	259
10.1.2.1	Basic configurations.....	260
10.1.2.2	Creating an access SAP egress policy.....	260
10.1.2.3	Editing QoS policies.....	261
10.2	Service egress policy command reference.....	261
10.2.1	Command hierarchies.....	261
10.2.1.1	Configuration commands.....	261
10.2.1.2	Copy commands.....	262
10.2.1.3	Show commands.....	262
10.2.2	Command description.....	262
10.2.2.1	Configuration commands.....	262
10.2.2.2	Operational commands.....	274
10.2.2.3	Show commands.....	276

11	Schedulers.....	279
11.1	Overview.....	279
12	Slope QoS policies.....	282
12.1	Overview of buffer pools and slope policies.....	282
12.2	Basic configurations.....	282
12.2.1	Creating a slope QoS policy.....	282
12.2.1.1	Default slope policy values for the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	283
12.3	Applying slope policies.....	285
12.4	Deleting QoS policies.....	285
12.4.1	Removing a policy from the QoS configuration.....	286
12.5	Copying and overwriting QoS policies.....	286
12.6	Editing QoS policies.....	287
12.7	Slope QoS policy command reference.....	287
12.7.1	Command hierarchies.....	287
12.7.1.1	Configuration commands.....	287
12.7.1.2	Operational commands.....	288
12.7.1.3	Show commands.....	288
12.7.2	Command descriptions.....	288
12.7.2.1	Configuration commands.....	288
13	Remark policies.....	299
13.1	Overview of remark policies for 7210 SAS-K 2F1C2T.....	299
13.1.1	Configuration guidelines for 7210 SAS-K 2F1C2T.....	300
13.2	Overview of remark policies for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	300
13.2.1	Configuration guidelines for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	303
13.3	Basic configurations.....	304
13.3.1	Creating a remark policy.....	304
13.3.2	Editing QoS policies.....	305
13.4	Remark policy command reference.....	305
13.4.1	Command hierarchies.....	305
13.4.1.1	Configuration commands for 7210 SAS-K 2F1C2T.....	305
13.4.1.2	Configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	306
13.4.1.3	Operational commands.....	306

13.4.1.4	Show commands.....	306
13.4.2	Command descriptions.....	306
13.4.2.1	Configuration commands.....	307
13.4.2.2	Operational commands.....	317
13.4.2.3	Show commands.....	317
14	Standards and protocol support.....	320
14.1	BGP.....	320
14.2	Certificate Management.....	322
14.3	Ethernet.....	322
14.4	EVPN.....	323
14.5	Fast Reroute.....	324
14.6	Internet Protocol (IP) — General.....	324
14.7	IP — Multicast.....	327
14.8	IP — Version 4.....	328
14.9	IP — Version 6.....	329
14.10	IPsec.....	330
14.11	IS-IS.....	331
14.12	Management.....	332
14.13	MPLS — General.....	335
14.14	MPLS — GMPLS.....	335
14.15	MPLS — LDP.....	336
14.16	MPLS — MPLS-TP.....	336
14.17	MPLS — OAM.....	337
14.18	MPLS — RSVP-TE.....	337
14.19	OSPF.....	337
14.20	Pseudowire.....	339
14.21	Quality of Service.....	339
14.22	RIP.....	340
14.23	Timing.....	340
14.24	VPLS.....	341

List of tables

Table 1: Supported modes of operation and configuration methods.....

18

Table 2: Supported port modes by mode of operation.....

20

Table 3: 7210 SAS platforms supporting port modes.....

20

Table 4: Configuration process.....

21

Table 5: QoS policy types and descriptions for 7210 SAS-K 2F1C2T.....

32

Table 6: QoS policy types and descriptions for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....

33

Table 7: Default network QoS policy egress marking.....

37

Table 8: Default network policy dot1p mapping to FC.....

37

Table 9: Default network QoS policy MPLS LSP EXP classification to FC mapping.....

38

Table 10: Default network QoS policy DSCP classification to FC mapping.....

38

Table 11: Default network QoS policy dot1p to FC mapping for network egress.....

40

Table 12: Default network QoS policy dot1p to FC mapping for network ingress.....

40

Table 13: Default network queue policy definitions.....

41

Table 14: Default service ingress policy ID 1 definitions.....

45

Table 15: Service ingress QoS policy IP match criteria.....

47

Table 16: Service ingress QoS policy IPv6 match criteria.....

47

Table 17: Service ingress QoS policy MAC match criteria.....

47

Table 18: Packet fields available for match in QoS classification policy and ACL policy.....

48

Table 19: MAC match Ethernet frame types.....

50

Table 20: MAC match criteria frame type dependencies.....

51

Table 21: Default service egress policy ID 1 definition.....

52

Table 22: FCs.....	57
Table 23: Default slope policy definitions.....	71
Table 24: Output fields: port.....	80
Table 25: Output fields: network.....	90
Table 26: Output fields: network queue.....	91
Table 27: Output fields: SAP-ingress QoS policy.....	93
Table 28: Platforms supported for classification policies.....	94
Table 29: Output fields: dot1p classification.....	106
Table 30: Output fields: DSCP classification.....	108
Table 31: Output fields: MPLS LSP EXP classification.....	110
Table 32: Network policy ID #1 defaults for access-uplink ports.....	115
Table 33: Default network QoS policy ID #1 for dot1p to FC mapping for access-uplink ports.....	116
Table 34: DSCP and dot1p marking for 7210 SAS-K 2F1C2T.....	130
Table 35: DSCP and dot1p marking for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	131
Table 36: Default DSCP mapping table.....	133
Table 37: Output fields: QoS network.....	162
Table 38: Output fields: SGT-QoS application.....	167
Table 39: Output fields: SGT-QoS DSCP map.....	170
Table 40: Output fields: network queue.....	193
Table 41: SAP-ingress policy defaults.....	198
Table 42: Output fields: SAP ingress.....	251
Table 43: Output fields: dot1p classification.....	254
Table 44: Output fields: DSCP classification.....	256

Table 45: Output fields: show SAP egress..... 277

Table 46: Default slope policy values..... 284

Table 47: Output fields: QoS slope policy.....298

Table 48: Summary of remark policy and attachment points for 7210 SAS-K 2F1C2T..... 299

Table 49: Summary of remark policy and attachment points for 7210 SAS-K 2F6C4T and 7210 SAS-
K 3SFP+ 8C..... 301

Table 50: Output fields: remark policy..... 319

List of figures

Figure 1: Network and service traffic types and QoS model.....

35

Figure 2: Traffic queuing model for FCs.....

45

Figure 3: RED slope characteristics.....

70

Figure 4: Scheduler implementation.....

280

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 provided by 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-K 2F1C2T
- 7210 SAS-K 2F6C4T
- 7210 SAS-K 3SFP+ 8C

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 the following 7210 SAS platforms operating implicitly in the specified modes. See [Table 1: Supported modes of operation and configuration methods](#) for more information.

- **network mode of operation**
7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
- **access-uplink mode of operation**
7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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 26.x.Rx content and may include some content that will be released in later maintenance loads. Refer to the *7210 SAS Software Release Notes 26.x.Rx*, part number 3HE 27852 000x TQZZA, for information about features supported in each load of the Release 26.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 are 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 phrases "mode of operation" and "operating mode" refer 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 26.x.Rx*, part number 3HE 27852 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 are available; refer to 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. Refer to 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-K 2F1C2T		Implicit	Implicit		
7210 SAS-K 2F6C4T ¹	Port Mode Configuration ²	Port Mode Configuration ²	Implicit		

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

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

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.

³ Implicitly operates in [network](#) mode when [standalone](#) mode of operation is configured

⁴ Supports MPLS uplinks only and implicitly operates in network mode

- **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.

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

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

Platform	Port mode			
	Access	Network	Hybrid	Access-uplink
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

1.4 7210 SAS QoS configuration process

The following table lists the tasks necessary to configure and apply QoS policies. The *7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Quality of Service 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	DEI-based classification and marking
	• egress rate	Port-level egress rate limiting
	• accounting mode	Frame-based accounting
	• DSCP, dot1p, and MPLS EXP Classification	DSCP, dot1p, and MPLS EXP classification policies

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.

Area	Task	Chapter
	• network	Network QoS policies
	• network queue	Network queue QoS policies
	• service ingress	Service ingress QoS policies
	• service egress	Service egress QoS policies
	• schedulers	Schedulers
	• slope policies	Slope QoS policies
	• remark policies	Remark 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 QoS policies and mechanisms to classify, queue, shape, and mark traffic.

2.1 QoS policies overview

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



Note:

The QoS capabilities supported on different 7210 SAS platforms are different. That is, not all the platforms support all of the capabilities. Please read through the following chapters to know what is available 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 (for example, QinQ tunnel, dot1q tunnel, IP/MPLS tunnel, and so on) 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 appear as a Layer 2 path to the service data; however, the data is really traversing an QinQ or IP or IP/MPLS core. The tunnel from one edge device to the other edge device is provisioned with an encapsulation, and the services are mapped to the tunnel that most appropriately supports the service needs.

The 7210 SAS supports the following FCs, internally named: Network-Control, High-1, Expedited, High-2, Low-1, Assured, Low-2, and Best-Effort. See [FCs](#) for more information about 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, however, there are a few instances where the default QoS policy uses a different ID. The default QoS policy is used if no policy is explicitly applied.

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

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

2.1.1 Overview of QoS policies on 7210 SAS-K 2F1C2T

On 7210 SAS-K 2F1C2T, QoS policies are applied on service ingress, service egress and access uplink ports (ingress and egress) and define the following:

- classification rules to map traffic to FCs
- forwarding class association with queues
- queue parameters for shaping, scheduling, and buffer allocation
- option to associate FC with meters/policers and configure meter parameters for enforcing rate and control burst
- QoS marking in the packet header

There are the following types of QoS policies:

- service ingress policies for access SAP-ingress
- service egress policies for access SAP-egress
- network policies for access-uplink port ingress and egress
- network queue policies for access-uplink port egress
- slope policies for all queues
- remark policies for both access SAP-egress and access-uplink port egress
- dot1p and DSCP classification policies for access SAP-ingress and access-uplink port ingress

Service ingress QoS policies are applied to the customer-facing Service Access Points (SAPs) on access ports. Traffic that enters through the access SAP is classified to map it to a forwarding class (FC) and the user has an option to also assign a profile on SAP-ingress. An FC can be associated with either queues or policer/meter on service ingress. That is, service ingress FC can be configured to use queue or a meter allowing for some FCs to use queues and some others to use meters. When the FC is associated with a queue on access SAP-ingress (also known as service ingress), the profile determines the enqueueing priority for the packet, with in-profile packets having a higher chance of getting a buffer, than out-of-profile packets. The mapping of traffic to FC/queues can be based on combinations of customer QoS marking in the packet header (for example, IEEE 802.1p bits, IP DSCP bits, MAC address, and so on). The characteristics of the FC queues are defined within the policy with regard to the number of FC queues to use for unicast traffic and BUM (Broadcast, Unknown-unicast, and Multicast) traffic along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS). Each of the FCs can be associated with different parameters for unicast traffic and different parameters for multipoint (that is, BUM) traffic.

A service ingress QoS policy defines up to eight queues per policy, with up to two queues (that is, Unicast Queue Mapping and Multicast Queue Mapping) per FC. Unicast and multipoint traffic can be mapped to use the same queue or mapped to use different queues per FC with a maximum of up to two queues per FC, one each for unicast and for multicast traffic.

For VPLS, the following types of forwarding are supported:

- unicast
- multicast
- broadcast
- unknown

Multicast, broadcast, and unknown forwarding 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. Multicast, broadcast,

and unknown traffic types use the same multicast-queue mapping defined for FC. That is, a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined.

Service ingress policies provide an option to use a policer per FC, instead of a queue. It allows users to classify low-latency less bursty traffic on service ingress to an FC and use a policer to enforce rate. When an FC is associated with policer on access SAP-ingress, the profile determines the ingress color for the packet. It allows in-profile (green) packets to use the available tokens from the CIR rate first while out-of-profile packets use available tokens from the PIR rate first. This allows the user to prioritize in-profile packets over out-profile packets by ensuring that CIR rate is available for in-profile service ingress packets.

The mapping of traffic to FC meter can be based on combinations of customer QoS marking in the packet header (for example, IEEE 802.1p bits, IP DSCP bits, MAC address, and so on). The characteristics of the FC policer are defined within the policy with regard to the number of FC meters to use for unicast traffic and BUM traffic along with the policer rate (like CIR, PIR) and burst parameters (like CBS, MBS). Each FC can be associated with different parameters for unicast traffic and different parameters for multipoint (that is, BUM) traffic.

A service ingress QoS policy defines up to 16 meters per policy, with up to 2 meters per FC. Unicast and multipoint traffic can be mapped to use the same policer or mapped to use different policer per FC with a maximum of up to 2 policer per FC, one each for unicast and for multicast traffic. In the case of VPLS service, four 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 in a point-to-point fashion within the service. Multicast, broadcast, and unknown traffic types use the same multicast meter mapping defined for FC. That is, a separate meter for multicast, broadcast, and unknown unicast traffic types cannot be defined.

Service ingress QoS policy also provides an option to configure a mix of meters and queues per policy and per FC. That is, it is possible to use a queue for one of the FC used to forward bursty traffic while using a meter for another FC used to forward low-latency traffic. In addition, it is possible to configure a queue for unicast traffic type while using a meter for BUM traffic types or the other way around.

On service ingress when a combination of queues and meters are used, the option is available to configure the service ingress aggregate rate for queues and meters individually. That is, the service ingress aggregate policer rate enforces the rate across all the FC meters configured in the SAP while the service ingress aggregate shaper rate enforces the rate across all the FC queues configured in the SAP.

Service egress QoS policies are applied to SAPs and map FCs to service egress queues for a service. The system can allocate a maximum of eight queues per SAP for the eight FCs. All traffic types (that is, both unicast and BUM traffic types) share the same queue on service egress. A service egress QoS policy defines the FC queue characteristics and also defines how to remark the FC to priority bits in the packet header (for example, IEEE 802.1p bits in the Ethernet VLAN tag) in the customer traffic.

Network QoS policies are applied to access-uplink ports. Access-uplink ports are typically used to connect to the core network and forward customer traffic toward the core network. A network QoS policy defines both ingress and egress behavior. On access-uplink port ingress, traffic that enters through the access-uplink port is classified to map it to an FC and the user has an option to assign a profile. FC is associated with ingress queues on access uplink port ingress and the profile determines the enqueueing priority for the packet, with in-profile packets have a higher chance of getting the buffer, than out-of-profile packets. The mapping of traffic to FC queues is based on QoS marking (for example, IEEE 802.1p bits, IP DSCP bits).

The characteristics of the FC ingress queues are defined within the policy as to the number of FC queues for the unicast traffic type and BUM traffic types, along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS, and so on). Each of the FCs can be associated with different ingress and ingress queue parameters for unicast traffic type and for multipoint (that is, BUM) traffic type.

A network QoS policy defines up to eight ingress queues per policy, with up to two ingress queues per FC. Unicast and multipoint traffic can be defined to use the same queue or different ingress queues per FC.

For VPLS, the following types of forwarding are supported:

- unicast
- multicast
- broadcast
- unknown

Multicast, broadcast, and unknown types are sent to multiple destinations within the service while the unicast forwarding type is handled in a point-to-point fashion within the service. All these traffic types use the same queue (a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined).

On access-uplink port egress, the policy maps FC and profile state to dot1p or IP DSCP values for traffic to be transmitted out of the access-uplink port. All the access-uplink SAPs configured on the same access-uplink port use the same policy and the same set of FC queues. Traffic received and transmitted through all the access-uplink SAPs configured on a particular access-uplink port receive similar QoS treatment.

On egress, network queue policies are applied to access-uplink ports and map FCs to network-egress queues on access-uplink ports. The system allocates eight egress queues per access-uplink port for the eight FCs. The policy defines the FC queue characteristics (that is, CIR, PIR, CBS, MBS, and so on). All traffic types (that is, both unicast and BUM traffic types) share the same queue on access-uplink port egress. All the access-uplink SAPs configured on the same access-uplink port use the same policy and the same set of FC queues. Traffic transmitted through all the access-uplink SAPs configured on a specific access-uplink port receive similar QoS treatment.

Slope policies are applied to service ingress queues, service egress queues, access uplink port ingress queues, and access uplink port egress queues. Each of these queuing points allocates buffers from the buffer pool and implements WRED for congestion management. During congestion WRED is used to evaluate how buffers from the pool are allocated to different FCs and to in-profile and out-of-profile traffic within a specific FC. The slope policies define the WRED parameters to use for in-profile/high-priority packets and for out-of-profile/low-priority packets. The high-slope and low-slope define the parameters for in-profile/high-priority packets and for out-of-profile/low-priority packets respectively. In addition, the 7210 SAS-K 2F1C2T introduces the concept of ring and non-ring ports with an option for preferential allocation of traffic for ring ports. The system by default treats access-uplink ports as ring ports.

Remark policies are applied to access SAP-egress and access-uplink port egress. They are not directly associated with the SAP and access-uplink port egress. Instead they are associated with service egress policy and network QoS policy. They define the FC and profile to egress marking values (for example, IEEE 802.1p bits in the Ethernet VLAN tag) to use.

Dot1p classification and DSCP classification allows the user to define the map of dot1p bits and IP DSCP values to FC and assign the profile for the packet on access SAP-ingress and access-uplink port ingress.

One service ingress QoS policy and one service egress QoS policy can be applied to a SAP access. One network QoS can be applied to a specific access-uplink port. One network queue policy can be applied to the access uplink port. If no QoS policy is explicitly applied to a SAP or port, a default QoS policy is applied.

2.1.2 Overview of QoS policies on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, QoS policies are applied on service ingress, service egress and access uplink ports (ingress and egress) and define the following:

- classification rules to map traffic to FCs

- FC association with queues
- queue parameters for shaping, scheduling, and buffer allocation
- option to associate FC with meters/policers and configure meter parameters for enforcing rate and control burst
- QoS marking in the packet header

There are the following types of QoS policies:

- service ingress policies for access SAP-ingress
- service egress policies for access SAP-egress
- network policies for access-uplink port ingress and egress, network port ingress and egress, and hybrid port ingress and egress
- network queue policies for access-uplink port egress, network port egress, and hybrid port egress
- slope policies for all queues
- remark policies for both access SAP-egress, access-uplink port egress, network port egress, and hybrid port egress
- dot1p, IP DSCP, and MPLS EXP classification policies (for access SAP-ingress, access-uplink port ingress and network port ingress). Note that MPLS EXP classification policies are only for network ports

Service ingress QoS policies are applied to the customer-facing SAPs on access ports. SAPs that are configured on hybrid ports must use the service ingress policies for ingress classification to an FC and queue. Traffic that enters through the SAP is mapped to an FC, and the user has an option to also assign a profile on SAP-ingress. An FC can be associated with either queues or policers/meters on service ingress. The service ingress FC can be configured to use queues or a policer/meter, allowing for some FCs to use queues and some others to use policers/meters.

When the FC is associated with a queue on access SAP-ingress (also known as service ingress), the profile determines the enqueueing priority for the packet, with in-profile packets having a higher chance of getting a buffer, than out-of-profile packets. The mapping of traffic to FC/queues can be based on combinations of customer QoS marking in the packet header (for example, IEEE 802.1p bits, IP DSCP bits, MAC address, and so on). The characteristics of the FC queues are defined within the policy as to the number of FC queues to use for unicast traffic and BUM traffic along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS). Each of the FCs can be associated with different parameters for unicast traffic and different parameters for multipoint (that is, BUM) traffic.

A service ingress QoS policy defines up to 8 queues per policy, with up to two queues (that is, Unicast Queue Mapping and Multicast Queue Mapping) per FC. Unicast and multipoint traffic can be mapped to use the same queue or mapped to use different queues per FC with a maximum of up to two queues per FC, one each for unicast and for multicast traffic.

For VPLS, the following types of forwarding are supported:

- unicast
- multicast
- broadcast
- 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. Multicast, broadcast, and unknown traffic types use the same multicast-queue mapping defined for FC. A separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined.

Service ingress policies provide an option to use a policer per FC, instead of a queue. It allows users to classify low-latency less bursty traffic on service ingress to an FC and use a policer to enforce rate. When an FC is associated with policer on access SAP-ingress, the profile determines the ingress color for the packet. It allows in-profile (green) packets to use the available tokens from the CIR rate first while out-of-profile packets use available tokens from the PIR rate first. The user can prioritize in-profile packets over out-profile packets by ensuring that CIR rate is available for in-profile service ingress packets. The mapping of traffic to FC meter can be based on combinations of customer QoS marking in the packet header (for example, IEEE 802.1p bits, IP DSCP bits, MAC address, and so on).

The characteristics of the FC policer are defined within the policy as to the number of FC meter to use for unicast traffic and BUM traffic along with the policer rate (like CIR, PIR) and burst parameters (like CBS, MBS). Each of the FCs can be associated with different parameters for unicast traffic and different parameters for multipoint (that is, BUM) traffic. A service ingress QoS policy defines up to 16 meters per policy, with up to two meters per FC. Unicast and multipoint traffic can be mapped to use the same policer or mapped to use different policer per FC with a maximum of up to two policer per FC, one each for unicast and for multicast traffic.

For VPLS, the following types of forwarding are supported:

- unicast
- multicast
- broadcast
- 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. Multicast, broadcast, and unknown traffic types use the same multicast meter mapping defined for FC. A separate meter for multicast, broadcast, and unknown unicast traffic types cannot be defined.

Service ingress QoS policy also provides an option to configure a mix of meters and queues per policy and per FC. It is possible to use a queue for one of the FC used to forward bursty traffic while using a meter for another FC used to forward low-latency traffic. In addition, it is possible to configure a queue for unicast traffic type while using a meter for BUM traffic types or the other way around.

On service ingress, when a combination of queues and meters are used, the option is available to configure the service ingress aggregate rate for queues and meters individually. That is, the service ingress aggregate policer rate enforces the rate across all the FC meters configured in the SAP while the service ingress aggregate shaper rate enforces the rate across all the FC queues configured in the SAP.

Service egress QoS policies are applied to access SAPs and map FCs to service egress queues for a service. The system can allocate a maximum of eight queues per SAP for the eight FCs. All traffic types (that is, both unicast and BUM traffic types) share the same queue on service egress. A service egress QoS policy defines the FC queue characteristics and also defines how to remark the FC to priority bits in the packet header (for example, IEEE 802.1p bits in the Ethernet VLAN tag) in the customer traffic. SAPs configured on hybrid ports must use service egress policies for egress queuing and remarking.

Network QoS policies are applied to access uplink ports, network ports, and hybrid ports. For an overview of the network QoS policy applied to access-uplink ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms, see [Overview of network QoS policy applied to access-uplink ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#). For an overview of the network QoS policy applied to network ports and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms, see [Overview of network QoS policy applied to network and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#).

On egress, network queue policies are applied to access-uplink ports, network ports, and hybrid ports. Network queue policies map FCs to network egress queues on access-uplink ports, network ports, and

hybrid ports. The system allocates eight egress queues per port (per network, hybrid, or access-uplink port) for the eight FCs. The policy defines the FC queue characteristics (that is, CIR, PIR, CBS, MBS, and so on). All traffic types (that is, both unicast and BUM traffic types) share the same egress queue on access-uplink ports, network ports, and hybrid ports. All the access-uplink SAPs that are configured on the same access-uplink port use the same policy and the same set of FC queues. This means that traffic transmitted through all the access-uplink SAPs configured on a specific access-uplink port receive similar QoS treatment. All network IP interfaces configured on the same network port or hybrid port use the same policy and the same set of FC queues; traffic transmitted through all network IP interfaces configured on a network port or hybrid port receive similar QoS treatment.

Slope policies are applied to service ingress queues, service egress queues, access-uplink port ingress queues, access-uplink port egress queues, network port ingress queues, and network port egress queues. Each of these queuing points allocates buffers from the buffer pool and implements WRED for congestion management. During congestion WRED is used to evaluate how buffers from the pool are allocated to different FCs and to in-profile and out-of-profile traffic within a specific FC. The slope policies define the WRED parameters to use for in-profile/high-priority packets and for out-of-profile/low-priority packets. The high-slope and low-slope define the parameters for in-profile/high-priority packets and for out-of-profile/low-priority packets respectively. In addition, the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C introduce the concept of ring and non-ring ports with an option for preferential allocation of traffic for ring ports. The system treats network ports and access-uplink ports as ring ports by default.

Remark policies are applied to access SAP-egress, access-uplink port egress, network port egress, and hybrid port egress. They are associated with service egress policy and network qos policy. They define the FC and profile to egress marking values (for example, IEEE 802.1p bits in the Ethernet VLAN tag) to use.

The dot1p classification policy, IP DSCP classification policy, and MPLS EXP classification policy allow the user to define the map of dot1p bits, IP DSCP values, and MPLS EXP values, respectively, to FCs, and assign the profile for the packet. The dot1p classification policy and IP DSCP classification policy are available on access SAP-ingress, access-uplink port ingress, network port ingress, and hybrid port ingress. The MPLS EXP classification policy is available on network port ingress and hybrid port ingress.

One service ingress QoS policy and one service egress QoS policy can be applied to a specific access SAP. One network QoS can be applied to a specific access-uplink port. One network queue policy can be applied to the network port or hybrid port. If no QoS policy is explicitly applied to a SAP or port, a default QoS policy is applied.

2.1.2.1 Overview of network QoS policy applied to access-uplink ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Access-uplink ports are typically used to connect to the core network using QinQ or dot1q links and forward customer traffic toward the core network. A network QoS policy defines both ingress and egress behavior on a access-uplink port. On access-uplink port ingress, traffic that enters through the port is classified to map it to an FC and the user has an option to assign a profile.

FC is associated with ingress queues on access-uplink port ingress and the profile determines the enqueueing priority for the packet, with in-profile packets have a higher chance of getting the buffer, than out-of-profile packets. The mapping of traffic to FC ingress queues is based on QoS marking (for example, IEEE 802.1p bits, IP DSCP bits).

The characteristics of the FC ingress queues are defined within the policy with regard to the number of FC queues for unicast traffic type and BUM traffic type, along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS, and so on). Each of the FCs can be associated with different ingress queue parameters for unicast traffic type and for multipoint (that is BUM) traffic type. A network QoS policy defines

up to eight ingress queues per policy, with up to 2 ingress queues per FC. Unicast and multipoint traffic can be defined to use the same ingress queue or different ingress queues per FC.

For VPLS, the following types of forwarding are supported:

- unicast
- multicast
- broadcast
- unknown

Multicast, broadcast, and unknown types are sent to multiple destinations within the service, while the unicast forwarding type is handled in a point-to-point manner within the service. All these traffic types use the same queue (a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined).

On access-uplink port egress, the policy maps FC and profile state to any combination of dot1p and IP DSCP values for traffic to be transmitted out of the access-uplink port. All the access-uplink SAPs configured on the same access-uplink port use the same policy and the same set of FC queues. That is, traffic received and transmitted through all the access-uplink SAPs configured on a specific access-uplink port receive similar QoS treatment.

2.1.2.2 Overview of network QoS policy applied to network and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Network ports are typically used to connect to the core network using IP/MPLS tunnels and forward customer traffic toward the core network. A network QoS policy defines both ingress and egress behavior on a network port. On network port ingress, traffic that enters through the port is classified to map it to an FC and the user has an option to assign a profile.

FC is associated with ingress queues on network port ingress and the profile determines the en-queuing priority for the packet, with in-profile packets have a higher chance of getting the buffer, than out-of-profile packets. The mapping of traffic to FC ingress queues is based on QoS marking (for example, MPLS EXP bits, IEEE 802.1p bits, IP DSCP bits).

The characteristics of the FC ingress queues are defined within the policy as to the number of FC queues for unicast traffic type and BUM traffic type, along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS, and so on). Each of the FCs can be associated with different ingress and ingress queue parameters for unicast traffic type and for multipoint (that is BUM) traffic type. A network QoS policy defines up to eight ingress queues per policy, with up to two ingress queues per FC. Unicast and multipoint traffic can be defined to use the same ingress queue or different ingress queues per FC.

For VPLS, the following types of forwarding are supported:

- unicast
- multicast
- broadcast
- unknown

Multicast, broadcast, and unknown types are sent to multiple destinations within the service while the unicast forwarding type is handled in a point-to-point manner within the service. All these traffic types use the same queue (a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined).

On network port egress or hybrid port egress, the policy maps FC and profile state to any combination of MPLS EXP, dot1p, and IP DSCP values for traffic to be transmitted out of the network IP interface that is configured on the network port or hybrid port. All network IP interfaces configured on the same network port or hybrid port use the same policy and the same set of FC queues; traffic received and transmitted through all network IP interfaces configured on a specified network port or hybrid port receive similar QoS treatment.

2.1.3 Summary of major functions of QoS policies

The following tables list a summary of the major functions performed by QoS policies.

Table 5: QoS policy types and descriptions for 7210 SAS-K 2F1C2T

Policy type	Applied at...	Description
Service ingress	Access SAP-ingress	- option to define up to eight FC queues and queue parameters to define queue characteristics (For example, scheduler priority and weight, rates, and so on) - option to defines up to 16 FC meters and meter parameters to define meter characteristics (For example, rates, CBS, MBS, and so on) - for traffic classification, defines match criteria to map flows to the queues based on dot1p or IP DSCP criteria or MAC criteria or IP criteria
Service egress	Access SAP-egress	- allocates up to eight FC queues and maps FCs to the queues. - defines FC-to-remarking values, through the use of remark policies - defines queue parameters to define queue characteristics (For example, scheduler priority and weight, rates, and so on)
Egress rate	Access port and Access-uplink port	configures the maximum bandwidth available for traffic sent out of a specified port
Network QoS	Access-uplink port	- at ingress, defines up to eight FC queues and queue parameters to define queue characteristics (For example, scheduler priority and weight, rates, and so on) - for traffic classification, defines match criteria to map flows to the queues based on dot1p and DSCP values - at egress, defines FC to remarking values, through the use of remark policies
Network queue	Access-uplink port	- allocates up to eight FC queues and maps FCs to the queues. - defines queue parameters to define queue characteristics (For example, scheduler priority and weight, rates, and so on)
Slope policies	SAP queues (both ingress)	enables or disables the high-slope and low-slope parameters for queue

Policy type	Applied at...	Description
	and egress) and Access-uplink port (both ingress and egress queues)	in addition to high-slope and low-slope, user has an option to use high-slope-ring and low-slope-ring parameters for access-uplink port egress queues
Remark policies	SAP-egress, access-uplink port egress	defines FC-to-remarking values; Not directly associated with a SAP or a port. Instead it is associated with SAP-egress policy and network qos policy
dot1p classification policy and DSCP classification policy	Access SAP-ingress and access-uplink port ingress	defines the map of dot1p bits and IP DSCP values to FC and assign the profile for the packet on access SAP-ingress and access-uplink port ingress

Table 6: QoS policy types and descriptions for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Policy type	Applied at...	Description
Service Ingress	Access SAP-ingress (SAP can be configured on access ports or hybrid ports)	- option to define up to eight FC queues and queue parameters to define queue characteristics (For example, scheduler priority and weight, rates, and so on) - option to defines up to 16 FC meters and meter parameters to define meter characteristics (For example, rates, CBS, MBS, and so on) - for traffic classification, defines match criteria to map flows to the queues based on dot1p or IP DSCP criteria or MAC criteria or IP criteria
Service egress	Access SAP-egress (SAP can be configured on access ports or hybrid ports)	- allocates up to eight FC queues and maps FCs to the queues - defines FC-to-remarking values, through the use of remark policies - defines queue parameters to define queue characteristics (For example, scheduler priority and weight, rates, and so on)
Egress rate	Access port, access-uplink port, network port, and hybrid port	configures the maximum bandwidth available for traffic sent out of a specified port
Network QoS	Network port, hybrid port, and access-uplink port	- at ingress, defines up to eight FC queues and queue parameters to define queue characteristics (For example, scheduler priority and weight, rates, and so on) - for traffic classification, defines match criteria using packet header bits (e.g. MPLS EXP, dot1p, IP DSCP, and so on) to map flows to the queues


Policy type	Applied at...	Description
		at egress, defines FC to remarking values, through the use of remark policies
Network queue	Network port Hybrid port Access-uplink port	- allocates up to eight FC queues and maps FCs to the queues - defines queue parameters to define queue characteristics (For example, scheduler priority and weight, rates, and so on)
Slope policies	Access SAP queues (both ingress and egress), Network port and hybrid port queues (both ingress and egress), and access-uplink port (both ingress and egress queues)	- enables or disables the high-slope and low-slope parameters for queue - in addition to high-slope and low-slope, user has an option to use high-slope-ring and low-slope-ring parameters for network port and access-uplink port egress queues
Remark policies	Access SAP-egress, network port egress (SAP can be configured on access ports or hybrid ports), hybrid port egress, and access-uplink port egress	defines FC-to-remarking values
dot1p classification policy and DSCP classification policy	Access SAP-ingress, network port ingress, hybrid port ingress, and access-uplink port ingress	- defines the map of dot1p bits and IP DSCP values to FC and assign the profile for the packet received on access SAP-ingress and access-uplink port - defines the map of dot1p and IP DSCP values to FC and assign the profile for packets received on network port
MPLS EXP classification policy	Network port ingress and hybrid port ingress	defines the map of MPLS EXP values to FC and assign the profile for packets received on network port

2.2 Network and service QoS policies

The QoS mechanism within the 7210 SAS is specialized to cater to the type of traffic on the interface.

The following figure shows that for customer interfaces, there are service ingress and service egress traffic types, and for access-uplink ports and network ports, there are network ingress and network egress traffic types.

Figure 1: Network and service traffic types and QoS model



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

The 7210 SAS supports the following major types of network and service QoS policies: [Network QoS policies on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C network ports and hybrid ports](#), [Network QoS policies on access-uplink ports](#), [Network queue policies](#), [Service ingress QoS policies](#), and [Service egress QoS policies](#).

The support of different policies varies across different platforms.

2.2.1 Network QoS policies on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C network ports and hybrid ports

Network QoS policies define egress QoS marking and ingress QoS interpretation for traffic received on network ports and hybrid ports.

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

- ingress
 - option to use MPLS EXP value to map traffic to available FCs and profile state
 - option to use dot1p value to map traffic to available FCs and profile state
 - option to use IP DSCP value to map traffic to available FCs and profile state
 - option to use all of three above simultaneously along with DEI for FC determination and assigning profile
 - defines FC-to-queue mapping
- egress

- option to define the FC and profile state to MPLS EXP value markings
- option to define the FC and profile state to dot1p value markings
- option to define the FC and profile state to IP DSCP value marking
- remarking of QoS bits can be enabled or disabled

The required network QoS policy definitions are the following:

- a unique network QoS policy ID
- egress FC to priority bits (for example, 802.1p, and so on) used for marking, for each FC
- a default ingress FC and an optional in-profile/out-of-profile state
- at least one default unicast FC meter or queue based on the platform. See [Meter/policers and policer parameters](#) for information about the parameters that can be configured for a meter

The optional network QoS policy element definitions are the following:

- additional queues
- MPLS EXP value to FC and profile state mappings for all values received
- dot1p value to FC and profile state mappings for all values received
- option to use DEI bit along with dot1p classification for profile state mapping
- option to use IP DSCP value to FC and profile state mappings for all DSCP values received

2.2.1.1 Ingress classification support for network ports and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

On ingress of network and hybrid ports, the user has an option to use both MPLS EXP and dot1p bits to map received MPLS packets to an FC. If both MPLS EXP and dot1p bits are configured, then the match order for MPLS packets is to match with MPLS EXP entries first and assign an FC if there is match. If no match exists, MPLS packets are matched with configured dot1p entries and assigned an FC if there is a match. If there is no match with both MPLS EXP and dot1p entries, the default configured FC is assigned. The DEI bit can be used to assign profile state to the MPLS packets on network ingress.

On ingress of network and hybrid ports, the user has an option to use both IP DSCP and dot1p bits to map received IP packets (plain routed packets in the context of network IP interfaces configured on the network port or hybrid port) to FC. If both IP DSCP and dot1p bits are configured, then the match order for IP packets is to match with IP DSCP entries first and dot1p entries next. The FC and profile value configured in the entry that matches first is assigned to the packet. If there is no match with either IP DSCP or dot1p values, the default FC is assigned to the packet. The DEI bit can be used to assign a profile state to the IP packets on network ingress.

MPLS EXP classification entries that map MPLS EXP values to FC, dot1p classification entries that map dot1p bits to FC, and IP DSCP classification entries that map IP DSCP values to FC are defined using MPLS EXP classification policies, dot1p classification policies, and DSCP classification policies, respectively.

2.2.1.2 Egress marking support for network ports and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

On network port and hybrid port egress, the option to mark MPLS EXP and dot1p values for MPLS packets is provided. For IP packets sent out of network IP interfaces configured on a network port or hybrid port, the option to mark IP DSCP and dot1p values for IP packets is provided. Along with dot1p, the user has an option to mark the DEI bit for both MPLS and IP DSCP packets.

Network policy ID 2 is reserved for the default network QoS policy applied to network and hybrid ports. The default policy cannot be deleted or changed. The default network QoS policy is applied to all network and hybrid ports that do not have another network QoS policy explicitly assigned.

The network QoS policy applied at the network egress (that is, on a network port or hybrid 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 packets, out-of-profile packets are preferentially dropped over in-profile packets at congestion points in the network.

For network egress, traffic remarking in the network QoS policy can be enabled or disabled.

The following table lists the default mapping of FC to dot1p values for egress marking.

Table 7: Default network QoS policy egress marking

FC-ID	FC name	FC label	Diff Serv name	Egress dot1p marking		MPLS EXP values		IP DSCP values	
				In-profile	Out-profile	In-profile	Out-profile	In-profile	Out-profile
7	Network Control	nc	NC2	111 - 7	111 - 7	7	7	NC2	NC2
6	High-1	h1	NC1	110 - 6	110 - 6	6	6	NC1	NC1
5	Expedited	ef	EF	101 - 5	101 - 5	5	5	EF	EF
4	High-2	h2	AF4	100 - 4	100 - 4	4	4	AF41	AF41
3	Low-1	l1	AF2	011 - 3	010 - 2	3	3	AF21	AF22
2	Assured	af	AF1	011 - 3	010 - 2	2	2	AF11	AF12
1	Low-2	l2	CS1	001 - 1	001 - 1	1	1	CS1	CS1
0	Best Effort	be	BE	000 - 0	000 - 0	0	0	BE	BE

For network ingress, the following figure lists the default mapping of dot1p values to FC and profile state for the default network QoS policy.

Table 8: Default network policy dot1p mapping to FC

dot1p value	FC	Profile
0	be	Out

dot1p value	FC	Profile
1	l2	In
2	af	Out
3	af	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

For network ingress, the following table lists the default mapping of mpls-lsp-exp-classification values to FC and profile state for the default network QoS policy.

Table 9: Default network QoS policy MPLS LSP EXP classification to FC mapping

MPLS EXP value	Ingress FC	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

For network ingress, the following table lists the default mapping of dscp-classification values to FC and profile state for the default network QoS policy.

Table 10: Default network QoS policy DSCP classification to FC mapping

IP DSCP value	Ingress FC	Profile
be	be	Out
ef	ef	In
cs1	l2	In
nc1	h1	In

IP DSCP value	Ingress FC	Profile
nc2	nc	In
af11	af	In
af12	af	Out
af41	h2	In

2.2.2 Network QoS policies on access-uplink ports

Network QoS policies define egress QoS marking and ingress QoS interpretation for traffic on received on access-uplink ports.

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

- **ingress**
 - option to use dot1p value mapping to FCs and profile
 - option to use IP DSCP value to map traffic to different FCs and profile
 - defines FC to ingress queue mapping
- **egress**
 - option to define the FC and profile to dot1p value markings
 - option to define the FC and profile to IP DSCP value marking
 - remarking of QoS bits can be enabled or disabled

The required network QoS policy element definitions are the following:

- a unique network QoS policy ID
- egress - FC and optional profile state to priority bits (for example, 802.1p, and so on) used for marking, for each FC
- a default ingress FC and an optional in-profile/out-of-profile state
- at least one default unicast FC queue. The parameters that can be configured for a ingress queue are discussed below.

The optional network QoS policy element definitions are the following:

- additional queues
- dot1p value to FC and profile state mappings for all values received
- option to use DEI bit along with dot1p classification for profile state mapping
- option to use IP DSCP value to FC and profile state mappings for all DSCP values received

2.2.2.1 Ingress classification support for access-uplink ports

On ingress of access-uplink ports, users have the option to use both IP DSCP and dot1p bits to map received IP packets to FC. If both IP DSCP and dot1p bits are configured, then the match order for IP

packets is to match with IP DSCP entries first and dot1p entries next. The FC and profile value configured in the entry which matches first is assigned to the packet. If there is no match with either IP DSCP and dot1p values, then the default FC is assigned to the packet. DEI bit can be used to assign profile state to the packets of network ingress.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, dot1p classification entries that map dot1p bits to FC and IP DSCP classification entries that map IP DSCP values to FC are defined by using dot1p-classification policies and DSCP classification policies respectively.

Network policy ID 1 is reserved for the default network QoS policy applied to access-uplink ports. The default policy cannot be deleted or changed. The default network QoS policy is applied to all access-uplink ports which do not have another network QoS policy explicitly assigned.

The network QoS policy applied at network egress (that is, 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 packets, out-of-profile packets are preferentially dropped over in-profile packets at congestion points in the network.

For network egress, traffic remarking in the network QoS policy can be enabled or disabled.

For network egress, the following table lists the default mapping of dot1p values to FC and profile state for the default network QoS policy.

Table 11: Default network QoS policy dot1p to FC mapping for network egress

dot1p value	Ingress FC	Profile
0	be	Out
1	l2	In
2	af	Out
2	l1	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

For network ingress, the following table lists the default mapping of dot1p values to FC and profile state for the default network QoS policy.

Table 12: Default network QoS policy dot1p to FC mapping for network ingress

dot1p value	Ingress FC	Profile
0	be	Out
1	l2	In
2	af	Out

dot1p value	Ingress FC	Profile
3	af	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

2.2.3 Network queue policies

Network queue policies are applied on egress of access-uplink ports on the 7210 SAS-K 2F1C2T.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, network queue policies are applied on egress for access-uplink ports, network ports, and hybrid ports.



Note:

The network egress aggregate shaper rate can be configured for hybrid ports on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C using the **nw-egr-agg-shaper-rate** command. Configuring the command limits the amount of traffic sent out of network IP interfaces configured on the hybrid port.

The network queue policies can be defined with up to a maximum of eight egress queues. The user has an option to define the policies with fewer than eight egress queues.

Queue characteristics configured on a per-FC basis are the following:

- committed buffer size (CBS) in Kilobytes
- maximum buffer size (MBS) in Kilobytes
- peak information rate (PIR) as a percentage of egress port bandwidth
- committed information rate (CIR) as a percentage of egress port bandwidth
- queue priority and queue weight

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

Table 13: Default network queue policy definitions

FC	Queue	Definition
Network-Control (nc)	Queue 8	PIR = 100% CIR = 10% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F1C2T

FC	Queue	Definition
		CBS = 24 Kilobytes -7210 SAS-K 2F6C4T CBS = 24 Kilobytes -7210 SAS-K 3SFP+ 8C priority = 1 weight = 1
High-1 (h1)	Queue 7	PIR = 100% CIR = 10% - 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T CIR = 5% - 7210 SAS-K 3SFP+ 8C MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F1C2T CBS = 24 Kilobytes -7210 SAS-K 2F6C4T CBS = 24 Kilobytes -7210 SAS-K 3SFP+ 8C priority = 1 weight = 1
Expedited (ef)	Queue 6	PIR = 100% CIR = 100% - 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T CIR = 15% - 7210 SAS-K 3SFP+ 8C MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F1C2T CBS = 24 Kilobytes -7210 SAS-K 2F6C4T CBS = 24 Kilobytes -7210 SAS-K 3SFP+ 8C priority = 1 weight = 1
High-2 (h2)	Queue 5	PIR = 100% CIR = 100% - 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T CIR = 15% - 7210 SAS-K 3SFP+ 8C

FC	Queue	Definition
		MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F1C2T CBS = 24 Kilobytes -7210 SAS-K 2F6C4T CBS = 24 Kilobytes -7210 SAS-K 3SFP+ 8C priority = 1 weight = 1
Low-1 (l1)	Queue 4	PIR = 100% CIR = 25% - 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T CIR = 10% - 7210 SAS-K 3SFP+ 8C MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F1C2T CBS = 24 Kilobytes -7210 SAS-K 2F6C4T CBS = 24 Kilobytes -7210 SAS-K 3SFP+ 8C priority = 1 weight = 1
Assured (af)	Queue 3	PIR = 100% CIR = 25% - 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T CIR = 10% - 7210 SAS-K 3SFP+ 8C MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F1C2T CBS = 24 Kilobytes -7210 SAS-K 2F6C4T CBS = 24 Kilobytes -7210 SAS-K 3SFP+ 8C priority = 1 weight = 1
Low-2 (l2)	Queue 2	PIR = 100% CIR = 25% - 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T

FC	Queue	Definition
		CIR = 10% - 7210 SAS-K 3SFP+ 8C MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F1C2T CBS = 24 Kilobytes -7210 SAS-K 2F6C4T CBS = 24 Kilobytes -7210 SAS-K 3SFP+ 8C priority = 1 weight = 1
Best-Effort (be)	Queue 1	PIR = 100% CIR = 0% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F1C2T CBS = 24 Kilobytes -7210 SAS-K 2F6C4T CBS = 24 Kilobytes -7210 SAS-K 3SFP+ 8C priority = 1 weight = 1

2.2.4 Service ingress QoS policies

Service ingress QoS policies define ingress service FC queues or meters and map traffic flows to FC on access SAP-ingress. A SAP can be configured on both access ports and hybrid ports. The service ingress QoS policy support for SAPs is the same for access ports and hybrid ports.



Note:

Not all 7210 SAS platforms support queues and meters on service ingress. The support varies across different platforms. See to subsequent chapters and sections for more information.

On 7210 SAS platforms, when a service ingress QoS policy is created, it always has one queue defined that cannot be deleted. The queue is used to queue all the traffic, both the unicast traffic and the multipoint traffic. These queues exist within the definition of the policy. The queues only get instantiated in hardware when the policy is applied to a SAP. Multipoint queues are instantiated only if the SAP-ingress policy defines a multipoint queue. By default, software does not allocate any multipoint queues.

In the simplest service ingress QoS policy, all traffic is handled as a single flow and mapped to a single queue, including all flooded traffic.

The required elements to define a service ingress QoS policy are the following:

- a unique service ingress QoS policy ID


- a QoS policy scope of template or exclusive
- at least one default FC queue. See [Queues and queue parameters](#) for more information about the parameters that can be configured for a queue.

Optional service ingress QoS policy elements for 7210 SAS platforms include the following:

- additional unicast queues or multicast queues, up to eight
- additional unicast meters or multicast meters, up to 16
- QoS policy match criteria to map packets to an FC

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

Figure 2: Traffic queuing model for FCs



2.2.4.1 Default service ingress policy

Service ingress QoS policy ID 1 is reserved for the default service ingress policy. The default policy cannot be deleted or changed. The default service ingress policy is implicitly applied to all SAPs that do not explicitly have another service ingress policy assigned. In the default policy, all traffic is mapped to the default FC which uses a queue by default. The following table lists the characteristics of the default policy.

Table 14: Default service ingress policy ID 1 definitions

Characteristic	Item	Definition
Queue	Queue 1	One queue defined for all unicast traffic and multicast traffic: • Forward Class: best-effort (be) • CIR = 0 • PIR = max • MBS = 60 Kilobytes • CBS = 10 Kilobytes

Characteristic	Item	Definition
		- Priority = 1 - Weight = 1
Flows	Default FC (be)	One flow defined for all traffic: All traffic mapped to best-effort (be)

2.2.5 Service ingress classification

Mapping flows to FCs is controlled by comparing each packet to the match criteria in the Service Ingress QoS policy applied to an access SAP. The ingress packet classification to FC is subject to a classification policy provisioned.

2.2.5.1 Service ingress classification

On access SAP-ingress, the user has an option to use either dot1p classification, IPv4 DSCP classification, IPv4 packet header fields, IPv6 packet header fields, or MAC packet header fields. The dot1p or DSCP classification rules to be used are defined in the dot1p and DSCP classification policy and associated with the SAP-ingress policy. The DSCP and dot1p classification policies can be configured in the same QoS policy. The IPv4 or IPv6 or MAC criteria can be configured in the SAP-ingress policy.

When packets are received on an access SAP, the following steps are used to determine the FC to assign to the packet.

1. Match IP criteria entries with the IP packet header fields in the packet. Assign the FC corresponding to the first entry which matches with IP packet header field values in the packet. If it is not an IP packet or if there is no match, go to next step.
2. Match MAC criteria entries with the MAC packet header fields in the packet. Assign the FC corresponding to the first entry which matches with MAC packet header field values in the packet. If there is no match, go to next step.
3. Match the IP DSCP value in the packet with the value configured in each of the IP DSCP entry defined in the DSCP classification policy. Assign the FC corresponding to the first entry which matches with IP DSCP value in the packet. If it is not an IP packet or if there is no match, go to next step.
4. Match the dot1p value in the packet (if available) with the value configured in each of the dot1p entry defined in the dot1p classification policy. Assign the FC corresponding to the first entry which matches with dot1p value in the packet. If there is no match, go to next step.
5. Assign the default FC.

2.2.5.2 Service ingress classification — IP and MAC packet fields

The IP and MAC match criteria can be very basic or quite detailed. IP and MAC match criteria are assembled from policy entries. An entry is identified by a unique, numerical entry ID. A single entry cannot contain more than one match value for each match criteria. Each match entry has an action which 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.

The following tables list the packets fields used for match-criteria used for access SAP-ingress classification.

Table 15: Service ingress QoS policy IP match criteria

IP criteria	Services applicable on 7210 SAS-K 2F1C2T	Services applicable on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
IP DSCP	Access SAPs in Epipe, VPLS, IES, RVPLS services	Access SAPs in Epipe, VPLS, IES, VPRN, RVPLS services
- IP source address and mask, IP destination address and mask, IP protocol, TCP/UDP source port and fragment field - TCP/UDP destination port	Access SAPs in Epipe, VPLS, IES, RVPLS services	Access SAPs in Epipe, VPLS, IES, VPRN, RVPLS services

Table 16: Service ingress QoS policy IPv6 match criteria

IPv6 criteria	Services applicable on 7210 SAS-K 2F1C2T	Services applicable on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
DSCP value, Destination IPv6 address and mask match, Destination port TCP/UDP port match, Source IPv6 address and mask match, Source port TCP/UDP port match	Access SAPs in Epipe, VPLS, RVPLS services	Access SAPs in Epipe, VPLS, RVPLS services

Table 17: Service ingress QoS policy MAC match criteria

MAC match criteria	Services applicable on 7210 SAS-K 2F1C2T	Services applicable on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
IEEE 802.1p/dot1p value/mask (for both inner and outer tag separately), Source MAC address/mask, Destination MAC address/mask, EtherType Value/Mask, outer VLAN, and inner VLAN tag value	Access SAPs in Epipe, VPLS, RVPLS services	Access SAPs in Epipe, VPLS, RVPLS services

The following table lists the Packet Fields available for match in QoS classification policy and ACL policy for different SAPs.

Table 18: Packet fields available for match in QoS classification policy and ACL policy

Ingress SAP type	Packet contents (only Ethernet-II frames)	MAC address match	Inner VLAN ID and dot1p match	Outer VLAN ID and dot1p match	Etype match	IPv4/IPv6 criteria match
NULL SAP	Null tag	Yes	No	No	Yes	Yes
	Priority tag (both VID and dot1p)	Yes	No	Yes	Yes	Yes
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No
dot1q SAP (includes dot1q explicit null SAP and dot1q Default SAP)	Null tag	Yes	No	No	Yes	Yes
	Priority tag (both VID and dot1p)	Yes	No	Yes	Yes	Yes
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No
dot1q SAP (includes dot1q SAP, dot1q range SAP)	Null tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Priority tag (both VID and dot1p)	Invalid	Invalid	Invalid	Invalid	Invalid
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No
QinQ SAP - 0.* SAP (matches only null and priority tag packets)	Null tag	Yes	No	No	Yes	Yes
	Priority tag (both VID and dot1p)	Yes	No	Yes	Yes	Yes
	Single tag	Invalid	Invalid	Invalid	Invalid	Invalid

Ingress SAP type	Packet contents (only Ethernet-II frames)	MAC address match	Inner VLAN ID and dot1p match	Outer VLAN ID and dot1p match	Etype match	IPv4/IPv6 criteria match
	Two tags	Invalid	Invalid	Invalid	Invalid	Invalid
	Three or more tags	Invalid	Invalid	Invalid	Invalid	Invalid
QinQ SAP (*.*, Default QinQ SAP)	Null tag	Yes	No	No	Yes	Yes
	Priority tag (both VID and dot1p)	Yes	No	Yes	Yes	Yes
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No
QinQ SAP (includes Q1.* SAP)	Null tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Priority tag (both VID and dot1p)	Invalid	Invalid	Invalid	Invalid	Invalid
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No
QinQ SAP (includes Q1.0 SAP)	Null tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Priority tag (both VID and dot1p)	Invalid	Invalid	Invalid	Invalid	Invalid
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags (inner tag is a priority tag)	Yes	Yes	Yes	Yes	Yes
	Two tags (inner tag is not a priority tag)	Invalid	Invalid	Invalid	Invalid	Invalid

Ingress SAP type	Packet contents (only Ethernet-II frames)	MAC address match	Inner VLAN ID and dot1p match	Outer VLAN ID and dot1p match	Etype match	IPv4/IPv6 criteria match
	Three or more tags	Yes	Yes	Yes	No	No
QinQ SAP (includes Q1.Q2 SAP)	Null tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Priority tag (both VID and dot1p)	Invalid	Invalid	Invalid	Invalid	Invalid
	Single tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Two tags (inner tag is a priority tag)	Invalid	Invalid	Invalid	Invalid	Invalid
	Two tags (inner tag is not a priority tag)	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No

The 7210 SAS does not support configuring of the frame-type match criteria and the default frame type configured is Ethernet - II; see the following table.

Table 19: 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.
802dot2-llc	IEEE 802.3 Ethernet frame with an 802.2 LLC header. Only the source MAC and destination MAC address are compared for match criteria.
802dot2-snap	IEEE 802.2 Ethernet frame with 802.2 SNAP header. Only the source MAC and destination MAC address 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).



Note:

The default frame type configured is Ethernet-II.

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

Table 20: MAC match criteria frame type dependencies

Frame format	Source MAC	Dest MAC	IEEE 802.1p value	Etype value
802.3	Yes	Yes	Yes	No
802dot2-llc	Yes	Yes	Yes	No
802dot2-snap	Yes	Yes	Yes	No
ethernet-II	Yes	Yes	Yes	Yes

2.2.6 Service egress QoS policies

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

- per-service egress sub-rate capabilities
- more granular, fairer scheduling per-service into the access network
- per-service statistics for forwarded and discarded service packets

The sub-rate capabilities and per-service scheduling control are required to make multiple services per physical port possible. With egress shaping, it is possible to support more than one service per port. It prevents traffic from single service 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. In the simplest service egress QoS policy, all FCs are treated like a single flow and mapped to a single queue.

Service egress QoS policies also define egress queues, shaping, scheduling, and remarking behavior for SAPs configured on hybrid ports. The QoS behavior for SAPs configured on hybrid ports is the same as for access SAPs configured on access ports.

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

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

Optional service egress QoS policy elements include the following:

- additional queues up to a total of eight separate queues. An FC queue is shared by unicast and multipoint (BUM) traffic types mapped to that FC.
- IEEE 802.1p priority value remarking based on FC
- option to use IP DSCP values for marking based on FC

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. More complex service

queuing models are supported in the router where each FC is associated with a dedicated queue. The FC 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. If the packet is received, the FC is marked in the tunnel transport encapsulation (for example, QinQ encapsulated packet).

2.2.6.1 Default service egress policy

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

Table 21: Default service egress policy ID 1 definition

Characteristic	Item	Definition
Queues	Queue 1	One queue defined for all traffic classes: • CIR = 0 • PIR = max (line rate) • MBS = 60 Kilobytes • CBS = 10 Kilobytes • Priority= 1 • Weight= 1
Flows	Default Action	One flow defined for all traffic classes: • All traffic mapped to queue 1 with no marking of IEEE 802.1p values

2.3 Meters/policers

This section provides information about meters/policers.

2.3.1 Meter/policers and policer parameters

This section describes the available meter parameters for meters used in service ingress policies. The terms "meters" and "policers" are used interchangeably in this document to refer to metering or policing.



Note:

Not all 7210 SAS platforms support meters for all the QoS policies. In addition, the meter parameters support available varies across different platforms. See the platform-specific QoS overview sections and the chapters following to know the support available on different platforms.

Meters are available with SAP-ingress policies associated with access SAP-ingress.

2.3.1.1 Meter ID

The meter ID is used to uniquely identify the meter/policer. The meter ID is only unique within the context of the QoS policy in which the meter is defined. It allows user to define multiple Meters with different characteristics and identify them while associating it with different FCs.

The user has the option to allocate up to eight meters and assign them a meter ID along with the option to configure the meter parameters that determine the meter characteristics.

2.3.1.2 Committed information rate

The committed information rate (CIR) for a meter is the long-term average rate at which traffic is considered conforming traffic or in-profile traffic. The higher the rate, the greater the expected throughput. The user will be able to burst above the CIR and up to the PIR for brief periods of time. The time and profile of the packet is determined based on the burst sizes. See [Committed burst size](#) and [Maximum burst size](#).

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 the hardware that determine the operational CIR for the meter. The user has limited 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 meters](#) for more information about the interpretation of the administrative CIR.



Note:

The in-profile and out-profile values assigned determine the meter behavior for the packet. See [Ingress profile assignment](#) for information.

2.3.1.3 Peak information rate

The peak information rate (PIR) defines the maximum rate at which packets are allowed to exit the meter. The PIR does not specify the maximum rate at which packets can enter the meter; this is controlled 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 rates in the hardware that determine the operational PIR for the meter. The user has limited 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 [Adaptation rule for meters](#) for more information about the interpretation of the administrative PIR.



Note:

The in-profile and out-profile values assigned determine the meter behavior for the packet. See [Ingress profile assignment](#) for information.

2.3.1.4 Adaptation rule for meters

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

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

- **minimum**

The system finds the hardware supported rate that is equal to or higher than the specified rate.

- **maximum**

The system finds the hardware supported rate that is equal to or lesser than the specified rate.

- **closest**

The system finds the hardware supported rate that is closest to the specified rate.

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



Note:

The 7210 SAS software considers the adaptation rules and the hardware values available to determine the admin PIR/CIR rates.

To illustrate (the example that follows is only for illustration of the use of adaptation rule and the values provided below does not list the actual values supported in hardware), how the adaptation rule constraints 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 Kb/s and 150 Kb/s, respectively. If the adaptation rule is minimum, the operational CIR and PIR values will be 128 Kb/s and 192 Kb/s respectively, as it is the native hardware rate greater than or equal to the administrative CIR and PIR values.

If the adaptation rule is maximum, the operational CIR and PIR values are 64 kb/s and 128 kb/s. If the adaptation rule is closest, the operational CIR and PIR values are 64 kb/s and 128 kb/s, respectively, as those are the closest matches for the administrative values that are even multiples of the 64 kb/s rate step.

2.3.1.5 Committed burst size

The committed burst size 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. The burst size configured by the user affects the rate step-size used by the system. The system uses the step size in a manner that both the burst-size and rate parameter constraints are met.

2.3.1.6 Maximum burst size

For Two Rate Three Color Marking (trTCM) policer using two-rate three-color marker) 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 is complying with PIR. For Single Rate Three Color Marking (srTCM) (policer using single rate three color marker), 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 marked as red are dropped.

The operational MBS set by the system is adapted from the user-configured value by using the minimum constraint. The burst size configured by the user affects the rate step-size used by the system. The system uses the step size in a manner that both the burst-size and rate parameter constraints are met.

2.3.1.7 Meter counters

The router maintains counters for meters within the system for the purpose of granular billing and accounting.

Each meter maintains the following counters:

- counter to count packets and octets forwarded in-profile
- counter to count packets and octets forwarded out-of-profile
- counter to count packets and octets dropped

Accounting record support available on each of the platforms is listed in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C System Management Guide* under the "Accounting Records/Logs" section.

2.3.1.8 Meter modes

The 7210 SAS supports the following meter modes:

- srTCM - Single Rate Three Color Marking
- trTCM - Two Rate Three Color Marking

For srTCM, the CBS and MBS Token buckets are replenished at a single rate, that is CIR; however, for trTCM CBS and MBS buckets are individually replenished at CIR and PIR rates, respectively. In trTCM1, the policing algorithm is implemented as defined in RFC 4115.

2.3.1.9 Meter platform support

For all platforms, a service meter can be provisioned on access SAP-ingress within service ingress QoS policies.

2.3.2 Ingress profile assignment

The meter can operate in profile mode (also called color-aware mode) and non-profile mode (also called color-blind mode). On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, SAP ingress meters operate in either profile mode or non-profile mode.

To enable profile mode or color-aware mode, the user must assign the initial ingress profile explicitly using the in/out profile commands, which can be performed using the classification entries or DEI. If the **use-dei** command is enabled, the in/out profile value assigned by the user is ignored (DEI takes priority). Nokia recommends using the default value of "undefined" for the ingress profile when DEI is enabled. To enable color-blind metering, the user must not assign an initial ingress profile value (and the default undefined is used). With both color-aware and color-blind metering, the final color is assigned by the meter associated

with the FC, based on the configured rates. The packet within CIR is assigned a final profile value of in-profile, and a packet that exceeds CIR and is within PIR is assigned a final profile value of out-profile. Anything above PIR is dropped.

The following functionality is implemented to support color-aware and color-blind metering:

- ingress profile assignment as part of ingress classification (initial profile value)
- on SAP-ingress, if the operator trusts the markings in the packet received from the customer device, the user has the following options:
 - DEI can be used to assign the initial profile value; use of DEI is enabled per FC.
 - Packet priority fields (for example, dot1p, IP DSCP, and so on) can be used for classification to an FC.
 - Packet priority fields (for example, dot1p, IP DSCP, and so on) can be used for classification to an FC and to assign the profile.
 - Profile is assigned using DEI overrides and initial profile value is assigned using explicit classification rules; however, Nokia recommends that you do not assign a profile explicitly when DEI use is enabled.
- on SAP-ingress, if the packet header priority value is not trusted, the user has the following options:
 - Packet fields (for example, mac-criteria and ip-criteria) can be used for classification to an FC and to assign the profile.
 - If no classification entry matches or if the matched classification entry does not explicitly assign the profile, these packets are assigned a value of undefined.
- ingress profile assignment as part of ingress policing or metering (lets call this the final profile value). The user has the following options:
 - If the initial profile value is assigned to a packet and if it is in-profile, it attempts to take tokens from CBS bucket. If available, its final profile value is set to in-profile. If not enough tokens are in the CBS bucket, check the MBS bucket. If sufficient tokens are available in the MBS bucket, the packet's final profile value is set to out-profile. If there are not enough tokens available in the MBS bucket, it is dropped; this is the color-aware mode of operation for in-profile packets.
 - If the initial profile value is assigned to a packet and if it is out-profile, it attempts to take tokens from the MBS bucket. If available, its final profile value is set to out-profile. If there are not enough tokens in the MBS bucket, it is dropped; this is the color-aware mode of operation for out-profile packets.
 - If the initial profile value is assigned to a packet and if it is undefined, it attempts to take tokens from the CBS bucket. If available, its final profile value is set to in-profile. If not available, check for enough tokens in the MBS bucket and if available its final profile is set to out-profile. If there are not enough tokens in the MBS bucket, it is dropped; this is the color-blind mode of operation.

2.3.3 QoS override

The QoS Override feature support on access SAPs 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.

The values are taken from the SAP-Ingress policy, when the meter parameter values are not overridden.

The configuration guidelines of QoS Override are:

- QoS override commands can be used only for the meters or policers defined in the SAP-ingress policy used with access SAPs.

- 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 supported for service egress QoS policies used with access SAPs.
- QoS override commands are not supported for ingress and egress QoS policies used with access-uplink ports.
- QoS override commands are not supported ingress and egress QoS policies used with network ports.

2.3.3.1 Configuring meter override parameters

Example

The following is a sample meter override parameter configuration output.

```
*A:dut-h>config>service>epipe>sap>ingress>meter-override>meter$ info detail
-----
mode trtcm2
adaptation-rule cir min pir max
cbs 1000 kbits
mbs 2000 kbits
rate cir 1000 pir 10000
-----
```

2.4 FCs

The 7210 SAS supports multiple FCs and class-based queuing, so the concept of FCs is common to all of the QoS policies. Each FC (also called Class of Service (CoS)) is important only in relation to the other FCs. A FC provides 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. [Table 22: FCs](#) describes the eight FCs supported by 7210 SAS.

The following table lists the default definitions for the FCs. The FC behavior, in terms of ingress marking interpretation and egress marking, can be changed by QoS policies.

Table 22: FCs

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.
5	Expedited	EF	EF	Intended for delay/jitter sensitive traffic.

FC-ID	FC name	FC designation	DiffServ name	Notes
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	

2.4.1 Forwarding-class to queue ID mapping

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

2.4.2 FC to queue ID mapping

The user has the option to define up to eight queues with an option to define the FC to queue mapping in service ingress policy, service egress policy, network qos policy and network queue policy.

2.5 Schedulers

The 7210 SAS supports Strict Priority and WFQ mode of scheduling or a mix of both.

On the 7210 SAS-K 2F1C2T, schedulers are used at SAP-ingress, SAP-egress, access-uplink port ingress and access-uplink port egress.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, schedulers are used at SAP-ingress, SAP-egress, network port ingress, network port egress, hybrid port ingress, hybrid port egress, access-uplink port ingress, and access-uplink port egress.

The scheduler uses 2 loops: the CIR loop and PIR loop, each with 4 priorities. The configured priority of the queue determines the service order of the queue in the CIR loop and the PIR loop. The scheduler first goes through the CIR loop, where it services all the queues which are operating at less than CIR rate according to their priority (that is, higher priority queues get services earlier than lower priority queues). It then goes through the PIR loop, where it services all the queues which are operating above the CIR rate (but less than PIR rate) according to their priority (that is, higher priority queues get services earlier than lower priority queues).

If there are multiple queues configured with the same priority, in the CIR loop the queues are scheduled using WFQ, with the configured weight (that is, pir-weight) of the queue used to determine the proportion of the available bandwidth that is provided to the queue. In the PIR loop, the queues are scheduled using WFQ, with the configured weight (that is, pir-weight) of the queue used to determine the proportion of the available bandwidth that is provided to the queue (using WFQ).

2.6 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, and ICMP. The packets destined for the CPU are classified internally and are put into the correct CPU 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.

2.7 Egress port rate limiting

This feature allows the user to limit the bandwidth available on the egress of the port to a value less than the maximum possible link bandwidth. On some platforms, it also allows the user to control the amount of burst sent out.

2.8 Queue management

This section provides information about QoS queue management.

2.8.1 Queues and queue parameters

This section describes the queue parameters provisioned for queues used in service ingress policy, service egress policy, access egress policy, network qos policy and network queue policy.



Note:

Not all 7210 SAS platforms support queues for all the above policies. In addition, the queue parameters support available varies across different platforms. See platform specific QoS overview sections above and the chapter following to know the support available on different platforms.

Queues are available on different platforms.

On the 7210 SAS-K 2F1C2T, queues are available with the following:

- SAP-ingress policies associated with access SAP-ingress
- SAP-egress policies associated with access SAP-egress
- Network queue policies associated with access-uplink port egress
- Network QoS policies associated with access-uplink port ingress

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, queues are available with the following:

- SAP-ingress policies associated with access SAP-ingress
- SAP-egress policies associated with access SAP-egress
- Network queue policies associated with access-uplink port egress
- Network QoS policies associated with access-uplink port ingress

- Network queue policies associated with network port and hybrid port egress
- Network QoS policies associated with network port and hybrid port ingress

2.8.1.1 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 within which the queue is defined. It allows user to define multiple queues with different characteristics and identify them while associating it with different FCs.

The user has an option to allocate up to 8 queues and assign them a queue ID along with the option to configure some of the queue parameters that determine the queue characteristics.

2.8.1.2 Committed information rate

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

- **Minimum bandwidth guarantees**

Queue's CIR setting provides the bandwidth that is provided to 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 a 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. The in-profile and out-of-profile state of queue is linked to scheduler prioritizing behavior as discussed below.

- **Scheduler queue priority metric**

The scheduler serving a group of queues prioritizes individual queues based on their current CIR and PIR states. Queues operating below their CIR are always served before those queues operating at or above their CIR. Also see information about schedulers to know the scheduler behavior on different 7210 SAS platforms.



Note:

The in-profile and out-profile state of the ingress queue determines the packet's final profile state based on the queue CIR, PIR values. The in-profile and out-profile state of the ingress queue also interacts with the scheduler mechanism and provides the minimum and maximum bandwidth guarantees. This is true only for ingress queues and not for egress queues. That is, the in-profile and out-profile state of the egress queue does not change the packets final profile state based on the queue CIR, PIR values. The in-profile and out-profile state of the egress queue only interacts with the scheduler mechanism and provides the minimum and maximum bandwidth guarantees. See [Ingress profile assignment](#) for more information.

When defining the CIR for a queue, the value specified is the administrative CIR for the queue. User has some control over how the administrative CIR is converted to an operational CIR should the hardware not support the exact CIR and PIR combination specified. The interpretation of the administrative CIR is discussed below in [Adaptation rule for queues](#).

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. An 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 should this behavior be required.

The CIR of the queue is configurable under the different QoS policies that provide the option to configure the queue parameters; for example, under service ingress and service egress queue policies, access port policies, and network queue policies.

2.8.1.3 Peak information rate

The peak information rate (PIR) defines the maximum rate at which packets are allowed to exit the queue. It does not specify the maximum rate at which packets may enter the queue; this is governed by the queue's ability to absorb bursts. 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's PIR, CIR and the relative priority and/or 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 should the hardware not support the exact CIR and PIR values specified. The interpretation of the administrative PIR is discussed in [Adaptation rule for queues](#)

The PIR of the queue is configurable under the different qos policies that provide the option to configure the queue parameters – example under service ingress and service egress queue policies, access port policies, network queue policies, and so on

2.8.1.4 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 the queue will be created on to derive the operational rates. The administrative CIR and PIR rates are translated to actual operational rates enforced by the hardware queue. The rule provides a constraint used when the exact rate is not available.

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 equal to or lesser than the specified rate.
- **closest**
Find the hardware supported rate that is closest to the specified rate.

Depending on the platform on which the queue is provisioned, the actual operational CIR and PIR settings used by the queue are dependent on the method the hardware uses to implement and represent the mechanisms that enforce 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.



Note:

The 7210 SAS software considers the adaptation rules and the hardware values available to determine the admin PIR/CIR rates.

To illustrate (the example that follows is only for illustration of the use of adaptation rule and the values provided below does not list the actual values supported in hardware), how the adaptation rule constraints 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 kb/s and 150 kb/s, respectively. If the adaptation rule is minimum, the operational CIR and PIR values will be 128 kb/s and 192 kb/s respectively, as it is the native hardware rate greater than or equal to the administrative CIR and PIR values. If the adaptation rule is maximum, the operational CIR and PIR values will be 64 kb/s and 128 kb/s. If the adaptation rule is closest, the operational CIR and PIR values will be 64 kb/s and 128 kb/s, respectively, as those are the closest matches for the administrative values that are even multiples of the 64 kb/s rate step.

2.8.1.5 Committed burst size

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

The CBS of the queue is configurable under the different QoS policies, if supported by the platform, that provide the option to configure the queue parameters – example under service ingress and service egress queue policies, network queue policies, and so on. The CBS for a queue is specified in kbytes.

The CBS for the queues is user-configurable. By default, software assigns a default value. It can be modified by the user as per their needs. The default value are specified in the command description.

2.8.1.6 Maximum burst size

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

The MBS of the queue is configurable under the different QoS policies, if supported by the platform, that provide the option to configure the queue parameters – example under service ingress and service egress queue policies, network queue policies, and so on. The MBS for a queue is specified in kbytes.

The MBS for the queues is user-configurable. By default, software assigns a default value. It can be modified by the user as per their needs. The default values are specified in the command description.

2.8.1.7 Ingress profile assignment

Queues can operate in two modes – profile mode and non-profile mode.

On the 7210 SAS-K 2F1C2T, SAP-ingress queues and access uplink port ingress queues operate in either profile mode or non-profile mode.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, SAP-ingress queues, network port and hybrid port ingress queues, and access uplink port ingress queues operate in either profile mode or non-profile mode.

In profile mode, the profile defined in the policy is used to determine the WRED slope to use for ingress queuing, with "profile in" packets using high-slope and "profile out" packets using low-slope. The ingress queue shaper does not change the profile value assigned to a packet that has a user assigned profile

value. That is, if a user assigns a profile value of green and the packet exceeds the CIR rate of the shaper, it is not changed to yellow. Similarly, packets assigned yellow are not changed by the shaper. The color assigned by the user is also subsequently used at the egress queuing point to determine the slope to use.

In non-profile mode, the profile is not specified by the user (and the node maps it to a value of "undefined". The low WRED slope is used at the ingress queuing point, as all packets received are considered to be the same as "profile out". The packet is then assigned the profile by the ingress queue shaper. It is assigned "in" profile value if it is within the CIR and assigned "out" profile value if it exceeds the CIR. It is dropped if it exceeds the PIR rate of ingress queue shaper (except for packets which are assigned a profile value of "undefined" on ingress and where the shaper assigns the profile based on CIR/PIR rate). The profile assigned by the ingress queue shaper is also subsequently used at the egress queuing point to determine the slope to use.

The user is provided with different options to assign the profile to the packet (for example, DEI based). It is always assigned on ingress of the packet into the node. When the profile is assigned at the ingress, it is used at subsequent queuing points in the system. That is, subsequent modules and queuing points in the system do not change the profile assigned to the packet on ingress. The profile assigned at ingress is also used to subsequently assign different marking/remarking values to in-profile and out-of-profile packets, if this meets user needs.

2.8.1.8 Weight and priority

A priority and weight can be assigned to the queue. The priority determines the service order of the queue when the scheduler schedules multiple queues configured on the same port. The queue weight determines the proportion of the available bandwidth that the scheduler allocates to a queue.

2.8.1.9 Queue counters

The router maintains counters for queues within the system for granular billing and accounting.

Each queue maintains the following counters:

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

The counters available vary among the 7210 SAS platform. Not all the counters are supported on all the platforms. Additionally there are restrictions on the number of counters that can be used simultaneously with a single queue. Some platforms can only count octets or packets and other can count both packets and octets. Counter (and corresponding accounting record) support available on each of the platform is listed in the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C System Management Guide in the Accounting Records/Logs section.

2.8.1.10 Queue platform support

The following support is available.

- **For all platforms**

Service queue is provisioned on access SAP-ingress and access SAP-egress service queues within service ingress QoS policies and service egress QoS policies, respectively.

- **For all platforms**

Access-uplink port ingress queues are defined within network QoS policies.

- **For all platforms**

Access-uplink port egress queues are defined within network queue policies.

- **On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C**

Network port and hybrid port ingress queues are defined within network QoS policies.

- **On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C**

Network port and hybrid port egress queues are defined within network queue policies.

2.8.2 Buffer pools

Buffer pools are used to manage the packet buffer memory resources used to store packets and absorb bursts received on a queue.

The total amount of available buffers (approximately 64 Mb on the 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T, and approximately 512 Mb on the 7210 SAS-K 3SFP+ 8C) is divided among the five buffer pools listed below. In addition, the following buffers are reserved for system internal use (such as multicast queues):

- CBS buffer pool
- ingress non-ring MBS pool
- egress non-ring MBS pool
- ingress ring MBS pool
- egress ring MBS pool

CBS buffer pool is used to allocate buffers toward CBS configured for ingress and egress queues on the node and some internal system queues.

The CBS pool allocation is as follows:

- On the 7210 SAS-K 2F1C2T, the CBS pool is used to allocate buffers toward CBS configured for ingress and egress queues on access SAPs, and ingress and egress queues on access-uplink ports.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the CBS pool is used to allocate buffers toward CBS configured for ingress and egress queues on access SAPs, ingress and egress queues on access-uplink ports, and ingress and egress queues on network ports and hybrid ports.

The MBS pool is divided into four pools as shown above: the ingress and egress non-ring MBS buffer pool and the ingress and egress ring MBS buffer pool. The MBS buffer pools can be over-subscribed.

The ingress and egress non-ring MBS buffer pool is used to allocate buffers toward the MBS configured for ingress queues and egress queues respectively on non-ring ports and non-ring service objects. The ingress and egress ring MBS buffer pool is used to allocate buffers toward the MBS configured for ingress queues and egress queues respectively on ring ports and ring service objects.

Ring ports and non-ring ports, and the corresponding ring and non-ring buffer pool on different platforms, are assigned as follows:

- On the 7210 SAS-K 2F1C2T, the access ports are designated as non-ring ports by default; this designation cannot be changed. The ingress non-ring MBS pool is used to allocate buffers toward all ingress queues configured on access SAPs. Similarly, the egress non-ring MBS pool is used to allocate buffers toward all egress queues configured on access SAPs.
- On the 7210 SAS-K 2F1C2T, all access-uplink ports are designated as ring ports and the ingress and egress ring MBS buffer pool is used to allocate buffers toward access-uplink port ingress and egress queues, respectively.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the access ports are designated as non-ring ports by default; this designation cannot be changed. The ingress non-ring MBS pool is used to allocate buffers toward all ingress queues configured on access SAPs. Similarly, the egress non-ring MBS pool is used to allocate buffers toward all egress queues configured on access SAPs.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, all network ports and access-uplink ports are designated as ring ports by default; this designation cannot be changed. The ingress and egress ring MBS buffer pool is used to allocate buffers toward network port and access-uplink port ingress queues and network port and access-uplink port egress queues, respectively.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, when a hybrid port is configured, the behavior is as follows:
 - SAPs configured on hybrid ports are designated as non-ring service objects.
 - Network port IP interfaces configured on hybrid ports are designated as ring service objects.

The amount of memory allocated toward these pools is software-defined and not user-configurable. The **show pools port-id system** command can be used to display total amount of buffers per pool and the amount of buffers in use per pool.

For the 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T, the values are:

```
A:dut-i# show pools 1/1/1 system

=====
Pool Information
=====
Port                : 1/1/1
Application         : System

MMU Total           : 65536 KB

MMU CBS             : 14336 KB          MMU CBS In Use      : 2240 KB

Ingress Ring        : 11776 KB          Ingress Ring In Use: 0 KB
Ingress NonRing     : 11776 KB          Ingress NonRing In*: 0 KB

Egress Ring         : 11776 KB          Egress Ring In Use : 0 KB
Egress NonRing      : 11776 KB          Egress NonRing In *: 0 KB

----- snipped -----
```

For the 7210 SAS-K 3SFP+ 8C, the values are:

```
A:dut-i# show pools 1/1/1 system

=====
Pool Information
=====
Port                : 1/1/1
Application         : System
```

```

MMU Total          : 524288 KB
MMU CBS            : 105472 KB      MMU CBS In Use      : 2240 KB
Ingress Ring       : 102400 KB      Ingress Ring In Use: 0 KB
Ingress NonRing    : 102400 KB      Ingress NonRing In*: 0 KB
Egress Ring        : 102400 KB      Egress Ring In Use : 0 KB
Egress NonRing     : 102400 KB      Egress NonRing In *: 0 KB
----- snipped -----

```

2.8.2.1 Ring and non-ring buffer pool

When the 7210 SAS-K 2F1C2T is deployed in a ring environment, the access-uplink ports are typically used to connect the node to ring. Similarly, on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, users will typically use the network ports to join the node into a ring. Therefore, these ports are designated as the ring ports. These ring ports carry traffic from the head-end toward the node (that is, the 7210 SAS), dropping traffic off to user/customer locations. Simultaneously, these ring ports carry traffic from the user/customer to the head-end. That is, traffic received from the user/customer is added to the ring and sent out toward the service edge, where services are terminated. The traffic in both these directions is typically admitted into the ring after being shaped and groomed. In the upstream direction (that is, in the direction of customer to service edge) the SLA is enforced at service ingress points (that is, typically access SAPs) and the traffic is shaped and groomed, similarly in the downstream direction (that is, in the direction of service edge to customer) it is done by the service edge device or the access aggregation device. That is, the downstream traffic should typically be allowed to pass through the intermediate nodes of the ring, without contention with and prioritized over the traffic that is received from the customer and being added into the ring by the nodes on the ring.

On the 7210 SAS-K 2F1C2T, the access-uplink ports are designated as ring ports and access ports are designated as non-ring ports. Traffic going from any access-uplink to another access-uplink port is identified as ring traffic. Traffic going from an access port to any access-uplink port, or traffic going from any access-uplink port to an access port (in egress), or traffic going from an access port to another access port is identified as non-ring traffic.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the network ports and access-uplink ports are designated as ring ports and access ports are designated as non-ring ports. Traffic going from any network port or access-uplink to another network port or access-uplink port is identified as ring traffic. Traffic going from an access port to any network port or access-uplink port, or traffic going from any network port or access-uplink port to an access port (in egress), or traffic going from an access port to another access port is identified as non-ring traffic.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the SAPs configured on hybrid ports are designated as non-ring service objects, and network port IP interfaces are designated as ring service objects. Traffic going from any network port IP interface on a hybrid port to another network port IP interface on a network port or a hybrid port, or the other way around, is identified as ring traffic. Traffic going from a SAP configured on a hybrid port to any network port, access-uplink port, or access port, or to another SAP on the hybrid port, or the other way around, is designated as non-ring traffic.

To ensure that the traffic received on ring ports is prioritized over traffic received on non-ring access ports, a separate ring MBS buffer pool (one each for ingress and egress) is provided for traffic received and sent out of ring ports. In addition, for ring ports and service objects, such as network port egress, hybrid port egress, and access-uplink egress (where shaped customer (access) traffic and ring traffic share the ring pool), two additional ring slopes (for a total of four configurable WRED slopes) are provided to prioritize

the ring traffic. Each egress queue on the network port, hybrid port, or access-uplink port supports four slopes per queue — ring high-slope, ring low-slope, non-ring high-slope, and non-ring low-slope (in the CLI command, the ring slopes are configured using the high-slope-ring and low-slope-ring, and the non-ring slopes are configured using the high-slope and low-slope). Ring high-slope and ring low-slope are used for in-profile and out-of-profile QoS profile ring traffic. Non-ring high-slope and low-slope are used for in-profile and out-of-profile non-ring traffic. Slope parameters (start-avg, max-avg, max-prob) of four slopes can be configured such that the ring traffic is prioritized over the non-ring traffic (that is, traffic being added onto the ring) in congestion scenarios.

A separate non-ring MBS buffer pool for traffic received and sent out of access ports along with two configurable WRED slopes is supported. Each queue on the access ports supports two slopes per queue — non-ring high-slope and non-ring low-slope. Non-ring high-slope and low-slope is used for in-profile and out-of-profile non-ring traffic. The non-ring buffer pool (one each for ingress and egress) is used to allocate buffers for non-ring traffic.

The usage of buffer pools for different traffic flows is as follows.

On the 7210 SAS-K 2F1C2T, the ring and non-ring buffer pools are used by the following traffic flows:

- Traffic received on access-uplink SAP and sent out of access-uplink SAP, uses the ring MBS buffer pool for MBS buffers on access-uplink port ingress and access-uplink port egress. In this case, ring high-slope is used for in-profile traffic and ring low-slope is used for out-of-profile traffic for both access-uplink ingress and access-uplink egress.
- Traffic received on access SAP and sent out of access-uplink SAP, uses the non-ring MBS buffer pool for MBS buffers on access SAP-ingress and uses the ring MBS buffer pool for MBS buffers on access-uplink SAP-egress. In this case, non-ring high slope and non-ring low slope is used on both access SAP-ingress and access-uplink egress.
- Traffic received on access-uplink SAP and sent out of another access SAP uses ring MBS buffer pool for MBS buffers on access-uplink SAP-ingress and the non-ring MBS buffer pool for MBS buffers on access SAP-egress. In this case, ring high-slope and ring low-slope is used on access-uplink ingress and non-ring high-slope and non-ring low-slope is used on access egress.
- Traffic received on access SAP and sent out of another access SAP uses the non-ring MBS pool for MBS buffers for both access SAP-ingress and access SAP-egress. In this case, non-ring high-slope and non-ring low-slope is used on both access ingress and access egress.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the ring and non-ring buffer pools are used by the following traffic flows:

- Traffic received on network port and sent out of network port, uses the ring MBS buffer pool for MBS buffers on network port ingress and network port egress. In this case, ring high-slope is used for in-profile traffic and ring low-slope is used for out-of-profile traffic for both network port ingress and port egress.
- Traffic received on access SAP and sent out of port, uses the non-ring MBS buffer pool for MBS buffers on access SAP-ingress and uses the ring MBS buffer pool for MBS buffers on network port egress. In this case, non-ring high slope and non-ring low slope is used on both access SAP-ingress and network port egress.
- Traffic received on network port and sent out of another access SAP uses ring MBS buffer pool for MBS buffers on network port ingress and the non-ring MBS buffer pool for MBS buffers on access SAP-egress. In this case, ring high-slope and ring low-slope is used on network port ingress and non-ring high-slope and non-ring low-slope is used on access egress.
- Traffic received on access-uplink SAP and sent out of access-uplink SAP, uses the ring MBS buffer pool for MBS buffers on access-uplink port ingress and access-uplink port egress. In this case, ring

high-slope is used for in-profile traffic and ring low-slope is used for out-of-profile traffic for both access-uplink ingress and access-uplink egress.

- Traffic received on access SAP and sent out of access-uplink SAP, uses the non-ring MBS buffer pool for MBS buffers on access SAP-ingress and uses the ring MBS buffer pool for MBS buffers on access-uplink SAP-egress. In this case, non-ring high slope and non-ring low slope is used on both access ingress and access-uplink egress.
- Traffic received on access-uplink SAP and sent out of another access SAP uses ring MBS buffer pool for MBS buffers on access-uplink SAP-ingress and the non-ring MBS buffer pool for MBS buffers on access SAP-egress. In this case, ring high-slope and ring low-slope is used on access-uplink ingress and non-ring high-slope and non-ring low-slope is used on access egress.
- Traffic received on access SAP and sent out of another access SAP uses the non-ring MBS pool for MBS buffers for both access SAP-ingress and access SAP-egress. In this case, non-ring high-slope and non-ring low-slope is used on both access ingress and access egress.
- Traffic received on a network IP interface on a hybrid port and sent out of a network port or another IP interface on a hybrid port uses the ring MBS buffer pool for MBS buffers on network or hybrid port ingress and network or hybrid port egress. In this case, ring high-slope is used for in-profile traffic, and ring low-slope is used for out-of-profile traffic for both network and hybrid port ingress and egress.
- Traffic received on a SAP on a hybrid port and sent out of a network port or network IP interface on a hybrid port uses the non-ring MBS buffer pool for MBS buffers on SAP-ingress and the ring MBS buffer pool for MBS buffers on network port egress. In this case, non-ring high-slope and non-ring low-slope are used on both access SAP-ingress and network port egress.
- Traffic received on a network IP interface on a hybrid port and sent out of a SAP configured on a hybrid port uses the ring MBS buffer pool for MBS buffers on network port ingress and the non-ring MBS buffer pool for MBS buffers on SAP-egress. In this case, ring high-slope and ring low-slope are used on network port ingress and non-ring high-slope and non-ring low-slope are used on SAP-egress.
- Traffic received on a SAP on a hybrid port and sent out of an access-uplink SAP uses the non-ring MBS buffer pool for MBS buffers on SAP-ingress and uses the ring MBS buffer pool for MBS buffers on access-uplink SAP-egress. In this case, non-ring high-slope and non-ring low-slope are used on both SAP-ingress and access-uplink egress.
- Traffic received on an access-uplink SAP and sent out of a SAP on a hybrid port uses the ring MBS buffer pool for MBS buffers on access-uplink SAP-ingress and the non-ring MBS buffer pool for MBS buffers on SAP-egress. In this case, ring high-slope and ring low-slope are used on access-uplink ingress and non-ring high-slope and non-ring low-slope are used on SAP-egress.
- Traffic received on a SAP on a hybrid port and sent out of another SAP on the hybrid port uses the non-ring MBS pool for MBS buffers for both SAP-ingress and access SAP-egress. In this case, non-ring high-slope and non-ring low-slope are used on both SAP-ingress and SAP-egress.

2.8.2.2 Configuration guidelines for CBS and MBS

- For configuring the CBS value, the following must be considered:
 - If jumbo frames need to be accommodated, then CBS must be set to at least a minimum of 10 kbytes. The default value set for the queue allows for jumbo frames. It is recommended to set the CBS to twice the amount of maximum frame size the queues is expected to carry.
 - CBS pool cannot be oversubscribed.
- For configuring the MBS value, the following must be considered:

- MBS value determines the maximum delay a packet can experience when using that queue. It should be set to a value such that the delay is acceptable.
- It is recommended to set the minimum value for MBS to be about four to five times the maximum size of the frame the queue is expected to carry to ensure better scheduling performance.

2.8.3 Slope policies

The available buffer space is partitioned into buffer pools as described in [Buffer pools](#). The buffers for a queue are allocated from the buffer pool. Slope policies define the RED slope characteristics.

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



Note:

If WRED is not configured, then taildrop is used.

2.8.4 RED slopes

This section provides information about the operation and configuration of RED slopes.

2.8.4.1 Operation and configuration of RED slopes

Each queue provides the following options:

- an option to use two slopes per queue on non-ring ports – a high-priority RED slope and a low-priority RED slope
- an option to use four slopes per queue on ring ports – a non-ring high-priority RED slope, a non-ring low-priority RED slope, a ring high-priority RED slope, and a ring 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. See [Buffer pools](#) for more information.

By default, the high-priority and low-priority slopes are disabled.

When a queue depth exceeds the queue CBS, packets received on that queue must contend with other queues exceeding their CBS for shared buffers. To resolve this contention, RED slopes are used to determine buffer availability on a packet-by- packet basis. A packet that is 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 are classified as low priority or out-of-profile are handled by this low-priority RED slope.

2.8.4.2 Simplified overview of RED


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

- The RED function keeps track of shared buffer utilization and shared buffer average utilization.
- At initialization, the utilization is zero and the average utilization is zero.

- When each packet is received, the current average utilization is plotted on the slope to determine the packet discard probability.
- A random number is generated associated with the packet and is compared to the discard probability.
- The lower the discard probability, the lower the chances that the random number is within the discard range.
- 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.
- A packet is discarded if the utilization variable is equal to the shared buffer size, or if the utilized CBS (actually in use by queues, not just defined by the CBS) is oversubscribed and has stolen buffers from the shared size, lowering the effective shared buffer size equal to the shared buffer utilization size.
- The new shared buffer average utilization is used as the shared buffer average utilization next time a packet probability is plotted on the RED slope.
- 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.

The following figure shows how a RED slope is a graph with an X (horizontal) and Y (vertical) axis. The X-axis plots the percentage of shared buffer average utilization, ranging 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.

Figure 3: RED slope characteristics



OSSG020

The following describes the sections shown in the preceding figure:

- Section A is (0, 0) to (start-avg, 0). This is the part of the slope where 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.
- 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.
- 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.
- 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 results 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 different queues (for example, access port queues) using the shared portion of the buffer pool, including disabling the RED slope.

2.8.4.3 Slope policy parameters

The elements required to define a slope policy are:

- a unique policy ID
- the high-slope (for in-profile packets) and low-slope (for out-of-profile packets) per queue; configurable parameters on each slope are start-avg, max-avg, and max-prob
- the ring and non-ring high and low slopes for access-uplink port egress
- on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the ring and non-ring high and low slopes for network port and hybrid port egress

The following table lists the default slope policy definitions.

Table 23: Default slope policy definitions

Parameter	Description
high-slope	start-avg 70 Percent max-avg 90 Percent max-prob 80 Percent
low-slope	start-avg 50 Percent max-avg 75 Percent max-prob 80 Percent
high-slope-ring	start-avg 70 Percent max-avg 90 Percent max-prob 80 Percent
low-slope-ring	start-avg 50 Percent max-avg 75 Percent max-prob 80 Percent

2.9 Preclassification on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the front-panel ports oversubscribe the capacity of the forwarding processor, with the exception of ports 1/1/1 and 1/1/2 on the 7210 SAS-K 3SFP+ 8C, which are not oversubscribed. A system-defined preclassification scheme is implemented (it is not user-configurable) to prioritize ingress packets for processing by the forwarding processor. It prioritizes packets based on dot1p and DSCP and identifies some of the untagged Layer 2 control protocols into high-priority and low-priority queues maintained on ingress per port. The forwarding processor processes the high-

priority queue across all the ports before servicing packets from the lower-priority queues. In addition, the network ports are favored over access ports by allocating more weight to the network ports during scheduling.

2.10 QoS policy entities

Services are configured using default QoS policies. Additional policies must be explicitly created and associated. For 7210 SAS-K platforms, the following policies are configured by default:

- one default service ingress QoS policy
- one default service egress QoS policy
- two default network ingress QoS policies (Network 1 and Network 2)
- one default network queue policy

Only one ingress QoS policy and one egress QoS policy can be applied to a SAP, access-uplink port, network port, or hybrid port.

When you create a new 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, and meters for ingress policies and queues for egress policies. By default, all FCs are mapped to Queue 1.

QoS policies can be applied to the following service types.

Epipe and VPLS:

- On the 7210 SAS-K 2F1C2T, SAP-ingress policies and SAP-egress policies are supported on an Epipe access service access point (SAP), VPLS access SAP, RVPLS SAP, and IES SAP.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, SAP-ingress policies and SAP-egress policies are supported on an Epipe service access point (SAP), VPLS access SAP, RVPLS SAP, IES access SAP, and VPRN access SAP.

QoS policies can be applied to the following entities:

- network QoS policy on access uplink port (all platforms)
- network queue policy (egress) on access uplink port (all platforms)
- network QoS policy on network port and hybrid port in network mode (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C only)
- network queue policy (egress) on network port and hybrid port in network mode (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C only)

2.10.1 Summary of QoS policy support for hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The following list describes an overview of QoS policy support on hybrid ports:

- Network queue policies are supported for queue configuration of egress queues on hybrid ports.
- Network QoS policies are supported for hybrid ports. The behavior is similar to the behavior for network ports, supporting per-port ingress classification and queuing and egress marking.
- SAP-ingress QoS policies are supported for SAPs configured on hybrid ports. The behavior is similar to the behavior for SAP-ingress on access ports.

- The network egress aggregate shaper rate can be configured for hybrid ports on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C using the **nw-egr-agg-shaper-rate** command. This command limits the amount of traffic sent out of network IP interfaces configured on the hybrid port. For information about configuring the egress aggregate shaper rate, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Interface Configuration Guide*.

2.11 Configuration notes

The following information describes QoS implementation guidelines and restrictions:

- Creating additional QoS policies is optional.
- Default policies are created for service ingress, service egress, access service egress, network, network queue, slope, remark, dot1p and DSCP classification, and port scheduler. (the policy types created varies across the platforms)
- Associating a service, access, or access uplink with a QoS policy other than the default policy is optional.

3 DEI-based classification and marking

This chapter 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 and marking

The DEI bit in the received packet can be used to assign the ingress profile for the packet. If DEI equals zero in the received packet, the packet is considered in-profile or green and if DEI equals one, the packet is considered out-of-profile or yellow. The use of the DEI bit for ingress classification can be enabled per FC. For a specific FC, if the DEI bit is used for the ingress profile assignment, the profile defined in the ingress classification entry is ignored. See [Ingress profile assignment](#) for information about the behavior when a profile is assigned to the packet on ingress.

On the 7210 SAS-K 2F1C2T, DEI-based classification is supported on access SAP ingress and access-uplink port ingress.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, DEI-based classification is supported on access SAP ingress, network port ingress, hybrid port ingress, and access-uplink port ingress.

3.2 DEI-based marking

The DEI bit can be used to mark the packet to carry the profile (which is assigned by an operator's trusted node on 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 following support is available for DEI-based marking:

- Option to mark DEI bits for access SAP egress on access ports is supported on all 7210 SAS platforms as described in this document.
- Option to mark DEI bits for port egress on access-uplink ports is supported on all 7210 SAS platforms as described in this document.
- Option to mark DEI bits for port egress on network ports and hybrid ports is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.
- by default, in-profile packets are marked with a DEI bit of 0 and out-of-profile packets are marked with a DEI bit of 1. The user has the option to mark all the packets belonging to an FC with the same DEI value, regardless of its profile, by using the **force-de-mark** command.
- DEI bits can be marked only if the remark policy of remark-type dot1p or dot1p-dscp is used.



Note:

See [Network QoS policy command reference](#), [Service egress policy command reference](#), and [Service SAP QoS policy command reference](#) for information about the CLI commands for DEI.

4 Port-level egress rate limiting

This chapter provides information to configure the port-level **egress-rate** command using the CLI.

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 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 to 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 the QoS capabilities and buffers required to prioritize the ingress traffic.

Egress rate attributes include the following:

- Per-port configuration of the maximum egress port rate is allowed, using the **egress-rate** CLI command.
- Ethernet ports configured as access and access-uplink support this feature.
- The port scheduler distributes the available maximum egress bandwidth based on the CIR/PIR configuration parameters provisioned for the queues.
- The burst parameter is not user-configurable and is set to a default by software.
- 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 (also known as frame-based accounting) is enabled, the queue scheduler accounts for the Ethernet frame overhead.
- When an egress-rate sub-rate value is provided, the access-uplink port egress queue rates that are specified using percentages will use the egress-rate value instead of the port bandwidth to configure the appropriate queue rates if the egress rate is less than the port bandwidth. Configuration of egress port rate to different values will result in a corresponding dynamic adjustment of rates for the egress queues configured on access-uplink ports.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, when an egress-rate sub-rate value is provided, the network or hybrid port egress queue rates that are specified using percentages use the egress-rate value instead of the port bandwidth to configure the queue rates if the egress rate is less than the port bandwidth. Manually configuring the egress port rate to different values results in a corresponding dynamic adjustment of rates for the egress queues that are configured on the network ports.

For hybrid ports, port queue rates specified as a percentage in the QoS policy change based on the lowest values of **nw-egr-aggr-shaper-rate**, **egress-rate**, or the port bandwidth.

- When the egress-rate sub-rate value is set, CBS/MBS of the associated network queues is not modified automatically. The user has an option to change the CBS/MBS values if necessary.

4.2 Basic configurations

For port-level rate limiting, the following considerations apply:

- The **egress-rate** command is in the **config>port>ethernet** context.
- The **egress-rate** command configures the maximum rate (in kb/s).

By default, the **egress-rate** command is not set on the port, and the port operates at the maximum line-rate speed it is operating at.

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
    exit
    no shutdown
-----
*A:Dut-1>config>port#
```

4.2.1 Modifying the port-level egress-rate command

To modify egress rate parameters, apply an **egress-rate** command with new **egress-rate** values.

4.2.2 Removing the port-level egress-rate command

To remove the **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
```

Example

The following is a sample configuration output that shows the removal of the **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#
```

4.3 Port level egress-rate command reference

- [Command hierarchies](#)
- [Command descriptions](#)

4.3.1 Command hierarchies

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

4.3.1.1 Configuration commands

```
- config
  - port
    - ethernet
      - egress-rate sub-rate
      - 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*

no egress-rate

Context

config>port>ethernet

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures maximum egress rate for a port. The egress-rate is configured as kb/s. The **no** form of this command removes egress-rate from the port.

Parameters

sub-rate
Specifies an integer value between 1 and 1000000 kb/s.

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

Description

This command displays egress rate and max burst value set for the port as well as other port details.

Parameters

port-id
Displays information about the specific port ID.

Output

The following output is an example of port information, and [Table 24: Output fields: port](#) describes the output fields.

Sample output

```
*A:dut-1>config>qos>network-queue# show port 1/1/1
=====
Ethernet Interface
=====
Description      : 10/100/Gig Ethernet SFP
Interface        : 1/1/1                Oper Speed      : 1 Gbps
Link-level       : Ethernet              Config Speed    : 1 Gbps
Admin State      : up                    Oper Duplex     : full
```

```

Oper State       : up                               Config Duplex   : full
Physical Link    : Yes                             MTU             : 1514
IfIndex          : 35684352                         Hold time up    : 0 seconds
Last State Change : 01/17/2011 04:05:37             Hold time down  : 0 seconds
Last Cleared Time : N/A

Configured Mode  : access                           Encap Type      : null
Dot1Q Ethertype  : 0x8100                          QinQ Ethertype  : 0x8100
Net. Egr. Queue Pol: default                       Access Egr. Qos *: 1
Egr. Sched. Pol  : default                         Network Qos Pol : n/a
Auto-negotiate   : limited                        MDI/MDX         : MDI
Accounting Policy : None                          Collect-stats    : Disabled
Egress Rate      : Default                        Max Burst       : Default
Uplink           : No

Down-when-looped : Disabled                       Keep-alive      : 10
Loop Detected    : False                         Retry           : 120

Configured Address : 00:78:76:45:54:02
Hardware Address   : 00:78:76:45:54:02
Cfg Alarm         :
Alarm Status      :

Transceiver Data

Transceiver Type  : SFP
Model Number      : 3HE00027AAA02 ALA IPUIAELDAB=
TX Laser Wavelength: 850 nm                      Diag Capable    : yes
Connector Code    : LC                          Vendor OUI       : 00:0a:1d
Manufacture date  : 2008/08/10                  Media           : Ethernet
Serial Number     : OPCPCH08052638
Part Number       : TRPAG1SXLAE5-TM
Optical Compliance : GIGE-SX
Link Length support: 550m for 50u MMF; 280m for 62.5u MMF;
=====
Traffic Statistics
=====
                                     Input      Output
-----
Octets                               0          0
Packets                             0          0
Errors                               0          0
=====
* indicates that the corresponding row element may have been truncated.
=====
Port Statistics
=====
                                     Input      Output
-----
Unicast Packets                      0          0
Multicast Packets                    0          0
Broadcast Packets                     0          0
Discards                             0          0
Unknown Proto Discards                0          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:dut-1>config>qos>network-queue#
```

Table 24: Output fields: port

Label	Description
Ethernet Interface	
Description	A text description of the port
Interface	The port ID displayed in the <i>slot/mda/port</i> format
Oper Speed	The operating 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	The operating duplex mode of the interface
Oper State	up — The port is operationally up down — The port is operationally down
Config Duplex	full — The link is configured to full-duplex mode half — The link is configured 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 that can be sent/received on the Ethernet physical interface, specified in octets
IfIndex	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 that reduces the number of link transitions reported to upper layer protocols.
Last State Change	The last time that the operational status of the port changed state
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

Label	Description
Configured Mode	network — The port is configured for transport network use access — The port is configured for service access hybrid — The port is configured for hybrid use (transport network and service access per VLAN)
Encap Type	null — Ingress frames do not use any tags or labels to delineate a service dot1q — Ingress frames carry 802.1Q tags, where each tag signifies a different service qinq — Ingress frames carry two 802.1Q tags, where the outer tag is the service provider tag and the inner tag is the customer service tag
Dot1Q Ethertype	The protocol carried in a dot1q Ethernet frame
QinQ Ethertype	The protocol carried in a QinQ Ethernet frame
Net.Egr. Queue Pol.	The number of the associated network egress queue QoS policy, or default if 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 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	The accounting policy applied to the port
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 will be 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 will not obtain the results and write them to the billing file.

Label	Description
Egress Rate	The maximum amount of egress bandwidth (in kilobits per second) that this Ethernet interface can generate
Down-when-looped	Enabled — The down-when-looped feature is enabled on the port Disabled — The down-when-looped feature is disabled on the port
Keep-alive	The time interval between keepalive PDUs transmitted toward the network during loop detection by the down-when-looped feature
Loop Detected	Indicates whether a loop is detected on the port
Retry	The minimum wait time before the port is re-enabled after it is brought down as a result of a loop detection
Configured Address	The base chassis Ethernet MAC address
Hardware Address	The interface hardware-assigned or system-assigned MAC address at its protocol sublayer
Cfg Alarm	The type of alarms to be logged and reported for the port
Alarm Status	The current alarm state
Transceiver Data	
Transceiver Type	The installed transceiver type
Model Number	The model number of the installed transceiver
TX Laser Wavelength	The wavelength of the transmission laser
Diag Capable	Displays whether digital diagnostic monitoring (DDM) is capable for the transceiver
Connector Code	The transceiver connector code
Vendor OUI	The vendor organizationally unique identifier (OUI)
Manufacture Date	The manufacture date of the transceiver
Media	The intended media for the transceiver to send and receive
Serial Number	The serial number of the transceiver
Part Number	The part number of the transceiver
Optical Compliance	The optical compliance code of the transceiver
Link Length Support	The supported link length of the transceiver

Label	Description
Traffic Statistics	
Octets input/output	The total number of octets received and transmitted on the port
Packets input/output	The number of packets, delivered by this sublayer to a higher (sub) layer, which were not addressed to a multicast or broadcast address at this sublayer. 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.
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. 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. For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors.
Port Statistics	
Unicast Packets Input/Output	The number of packets, delivered by this sublayer to a higher (sub) layer, which were not addressed to a multicast or broadcast address at this sublayer. 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.
Multicast Packets Input/Output	The number of packets, delivered by this sublayer to a higher (sub) layer, which were not addressed to a unicast or broadcast address at this sublayer. The total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a unicast or broadcast address at this sublayer, including those that were discarded or not sent
Broadcast Packets Input/Output	The number of packets, delivered by this sublayer to a higher (sub) layer, which were not addressed to a unicast or multicast address at this sublayer. The total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a unicast

Label	Description
	or multicast address at this sublayer, including those that were discarded or not sent.
Discards Input/Output	The number of inbound/outbound packets chosen to be discarded to possibly free up buffer space
Unknown Proto Discards Input/Output	For packet-oriented interfaces, the number of packets received via the interface that were discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing, the number of transmission units received via the interface that were discarded because of an unknown or unsupported protocol. For any interface that does not support protocol multiplexing, this counter will always be 0. Unknown proto discards do not show up in the packet counts
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
Sngl Collisions	The number of frames that are involved in a single collision, and are subsequently transmitted successfully
FCS Errors	The number of frames received that are an integral number of octets in length but do not pass the FCS check
Mult Collisions	The number of frames that are involved in more than one collision and are subsequently transmitted successfully
SQE Test Errors	The number of times that the SQE TEST ERROR is received
Late Collisions	The number of times that a collision is detected later than one slotTime into the transmission of a packet
CSE	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame
Excess Collisns	The number of frames for which a transmission fails as a result of excessive collisions
Too long Frames	The number of frames received that exceed the maximum permitted frame size

Label	Description
Int MAC Tx Errs	The number of frames for which a transmission fails as a result of an internal MAC sublayer transmit error
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
Int MAC Rx Errs	The number of frames for which a reception fails as a result of an internal MAC sublayer receive error

5 Frame-based accounting

This chapter provides information to configure frame-based accounting using the CLI.

5.1 Overview of frame-based accounting

When enabled, frame-based accounting allows QoS policies account for the Ethernet frame overhead (for example, it accounts for the IFG (inter-frame gap) and the preamble). Typically, the IFG and preamble constitute about 20 bytes (12 + 8). The QoS meter/policer and shaper use this overhead for Ethernet ports when allocating bandwidth.

On 7210 SAS platforms, a configurable CLI command enables accounting of the frame overhead per port. This command affects the behavior of the SAP ingress FC meter, SAP ingress aggregate meter, ingress queue rate, and egress queue rate of all the SAPs configured on the port. When disabled, the SAP ingress FC meter, SAP ingress aggregate meter, ingress queue rate, and egress queue rate, along with the port egress rate, do not account for the Ethernet frame overhead. When enabled, the SAP ingress FC meter, SAP ingress aggregate meter, ingress queue rate, and egress queue rate, along with the port egress rate, account for the Ethernet frame overhead. By default, frame-based accounting is disabled for the port.

Accounting records and statistics account for frame overhead for SAPs configured on the port when FBA is enabled on the port.

On the 7210 SAS-K 2F1C2T, frame-based accounting is supported on both access ports and access-uplink ports.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, frame-based accounting is supported on network ports, hybrid ports, access ports, and access-uplink ports.

5.2 Enabling and disabling frame-based accounting

On 7210 SAS platforms, frame-based accounting is supported per port with the capability to enable and disable it per port for both ingress and egress. In other words, it is not possible to enable or disable it only for ingress or for egress; both can be enabled together or disabled together.

To enable frame-based-accounting for both ingress and egress on a port, execute the command **config>port>ethernet>frame-based-accounting**.

To disable frame-based-accounting for both ingress and egress on a port, execute the command **config>port>ethernet>no frame-based-accounting**.

Example: Enabling frame-based accounting

```
*A:Dut-1>config>port>ethernet>#info detail
-----
...snipped...
      frame-based-accounting;
...snipped...
-----
*A:Dut-1>config>port>ethernet #
```

Example: Disabling frame-based accounting

```
*A:Dut-1>config>port>ethernet>#info detail
-----
...snipped...
      no frame-based-accounting;
... snipped ...
-----
*A:Dut-1>config>port>ethernet #
```

For more information about the command, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP + 8C Interface Configuration Guide*.

5.3 Frame-based accounting command reference

- [Command hierarchies](#)
- [Command descriptions](#)

5.3.1 Command hierarchies

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

5.3.1.1 Configuration commands

```
config
- port
  - ethernet
    - frame-based-accounting
    - no frame-based-accounting
```

5.3.1.2 Show commands

```
show
- qos
  - network [policy-id] [detail]
  - network-queue [network-queue-policy-name] [detail]
  - sap-ingress [policy-id] [association | match-criteria | detail]
```

5.3.2 Command descriptions

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

5.3.2.1 Configuration commands

frame-based-accounting

Syntax

frame-based-accounting

no frame-based-accounting

Context

config>port>ethernet

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configure per port frame-based accounting. It can be enabled or disabled on each port.

When enabled, all SAP ingress FC meter rates, SAP ingress aggregate meter rates, shaper rates, meter statistics, and queue statistics on that port also account for the Ethernet Layer 1 overhead (of 20 bytes) in both ingress and egress directions. For example, all SAP ingress FC meter rates, SAP ingress aggregate meter rates, ingress queue shaper rates, egress queue shaper rates, and aggregate SAP shaper rates account for the Ethernet overhead.

The **no** form of this command disables frame-based-accounting.

Default

no frame-based-accounting

5.3.2.2 Show commands

network

Syntax

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

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the accounting status of a network QoS policy along with other details of the policy. When frame-based-accounting is enabled accounting is shown as frame-based otherwise packet-based.

Parameters

- policy-id**
Displays information about the specific policy ID.
- detail**
Displays the detail policy information.

Output

The following output is an example of QoS network policy information, and [Table 25: Output fields: network](#) describes the output fields.

Sample output

```
*A:dut-a>show>qos# network 1

=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id      : 1
Egr Remark    : False                      Egr Rem Plcy   : N/A
Forward Class  : be                        Profile       : None
Scope         : Template
DOT1P Class Poli*: 1                      DSCP Class Polic*: 0
MPLS Lsp Exp Cla*: 0
Description    : Default network-port QoS policy.
=====
* indicates that the corresponding row element may have been truncated.
*A:dut-a>show>qos#
*A:dut-a>show>qos# network 1

=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id      : 1
Egr Remark    : False                      Egr Rem Plcy   : N/A
Forward Class  : be                        Profile       : None
Scope         : Template
DOT1P Class Poli*: 1                      DSCP Class Polic*: 0
MPLS Lsp Exp Cla*: 0
Description    : Default network-port QoS policy.
=====
* indicates that the corresponding row element may have been truncated.
*A:dut-a>show>qos#
```

Table 25: Output fields: network

Label	Description
Policy-ID	Displays the policy identifier
Profile	Out — Specifies the dot1p marking for the packets which are out-of-profile, egressing on this queue In — Specifies the dot1p markings for in-profile packets egressing this queue
Scope	Exclusive — Implies that this policy can only be applied to a single SAP Template — Implies that this policy can be applied to multiple SAPs on the router
Description	A text description of the port

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

Description

This command displays accounting status of a network-queue policy along with other details of the policy. When frame-based-accounting is enabled, accounting is shown as frame-based; otherwise it is shown as packet-based.

Parameters

network-queue-policy-name

Displays information about the specific Network queue policy.

detail

Displays the detailed policy information.

Output

The following output is an example of QoS network-queue policy information, and [Table 26: Output fields: network queue](#) describes the output fields.

Sample output

```
*A:Dut-1# show qos network-queue default
=====
QoS Network Queue Policy
=====
-----
Network Queue Policy (default)
-----
Policy          : default
Accounting       : frame-based
Description      : Default network queue QoS policy.
-----
Associations
-----
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 : 1/1/15
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:Dut-1#
```

Table 26: Output fields: network queue

Label	Description
Policy	Displays the policy
Accounting	Packet-based — Specifies that the meters associated with this policy do not account for packet framing overheads (such as the Inter Frame Gap (IFG) and the preamble for Ethernet), 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
Description	A text description of the port
Port-Id	Displays the specified port ID

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 accounting status of a sap-ingress policy along with other details of the policy. When frame-based-accounting is enabled accounting is shown as frame-based otherwise packet-based.

Parameters

- policy-id**
Displays information about the specific policy ID.
- associations**
Displays the associations of the sap-ingress policy.
- match-criteria**
Displays the match criteria of the sap-ingress policy.
- detail**
Displays the detailed information of the sap-ingress policy.

Output

The following output is an example of QoS SAP ingress policy information, and [Table 27: Output fields: SAP-ingress QoS policy](#) describes the output fields.

Sample output

```
*A:dut-a>show>qos# sap-ingress 1

=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (1)
-----
Policy-id           : 1           Scope           : Template
Default FC         : be
Criteria-type      : None
Mac Sub-Criteria   : None        IP Sub-Criteria   : None
IPv6 Enabled       : False
DOT1P Class Policy Id : 0         DSCP Class Policy Id : 0
MPLS Lsp Exp Class Policy*: 0
Name               : default
Description        : Default SAP ingress QoS policy.
```

```
=====
* indicates that the corresponding row element may have been truncated.
*A:dut-a>show>qos#
```

Table 27: 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
MAC Sub-Criteria	Displays the configured MAC sub-criteria
IP Sub-Criteria	Displays the configured IP sub-criteria
Description	A text string that helps identify the policy's context in the configuration file

6 DSCP, dot1p, and MPLS EXP classification policies

This chapter provides information to configure the DSCP classification policy, dot1p classification policy, and MPLS EXP classification policy using the CLI.

6.1 Overview

These policies allow the user to define a policy or template that maps the packet priority bits, like dot1p, IP DSCP, and MPLS EXP bits to FC and profile. The template can then be used in a SAP ingress policy to define the ingress classification of flows to FC.

The following table describes the support available on different 7210 SAS platforms.

Table 28: Platforms supported for classification policies

Policy name	Support available
dot1p-classification	All platforms with sap-ingress policy and network-qos policy.
DSCP classification	All platforms with sap-ingress policy and network-qos policy.
MPLS-EXP classification	Only 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C with network-qos policies

6.1.1 DSCP classification policy

This policy is used to define the map of IP DSCP values in the IP header of the packet to FC and profile (ingress profile).

6.1.2 dot1p classification policy

This policy is used to define the map of dot1p values in the Ethernet frame to FC and profile (ingress profile).

6.1.3 MPLS EXP classification policy

This policy is used to define the map of MPLS EXP values in the MPLS header of the packet to FC and profile (ingress profile).

6.2 Configuring classification policies

Example

```
config>qos>dot1p-classification 1 create
config>qos>dscp-classification 1 create
config>qos>mpls-lsp-exp-classification 1 create
```

Example

The following is a sample configuration output that shows the enabling of the DSCP classification policy, dot1p classification policy, and MPLS EXP classification policy.

```
*A:K-SASK12>config>qos># info detail
-----
-----exit
dot1p-classification 1 create
  description "Default Dot1P Classification policy"
  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
dot1p-classification 10 create
  no description
exit
dscp-classification 1 create
  description "Default DSCP Classification policy"
  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
dscp-classification 20 create
  no description
exit
mpls-lsp-exp-classification 1 create
  description "Default MplsLspExp Classification policy"
  lsp-exp 0 fc "be" profile out
  lsp-exp 1 fc "l2" profile in
  lsp-exp 2 fc "af" profile out
  lsp-exp 3 fc "af" profile in
  lsp-exp 4 fc "h2" profile in
  lsp-exp 5 fc "ef" profile in
  lsp-exp 6 fc "h1" profile in
  lsp-exp 7 fc "nc" profile in
exit
mpls-lsp-exp-classification 20 create
  no description
exit
```

```
*A:K-SASK12>config>qos>dot1p-classification#

*A:K-SASK12>config>qos>mpls-lsp-exp-classification# info detail
-----
      description "Default MplsLspExp Classification policy"
      lsp-exp 0 fc "be" profile out
      lsp-exp 1 fc "l2" profile in
      lsp-exp 2 fc "af" profile out
      lsp-exp 3 fc "af" profile in
      lsp-exp 4 fc "h2" profile in
      lsp-exp 5 fc "ef" profile in
      lsp-exp 6 fc "h1" profile in
      lsp-exp 7 fc "nc" profile in
-----
*A:K-SASK12>config>qos>mpls-lsp-exp-classification
```

6.3 DSCP, dot1p, and MPLS EXP classification policy command reference

- [Command hierarchies](#)
- [Command description](#)

6.3.1 Command hierarchies

- [Configuration commands for 7210 SAS-K 2F1C2T](#)
- [Configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [Show commands](#)
- [Operational commands](#)

6.3.1.1 Configuration commands for 7210 SAS-K 2F1C2T

```
- config
- qos
- dot1p-classification classification-id [create]
- no dot1p-classification classification-id
  - description description-string
  - no description
  - dot1p dot1p-priority fc fc-name profile {in | out}
  - no dot1p dot1p-priority
- dscp-classification classification-id [use-ip-prec-classification] [create]
- no dscp-classification classification-id
  - description description-string
  - no description
  - dscp dscp-name fc fc-name profile {in | out}
  - no dscp dscp-name
```

6.3.1.2 Configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
- config
```



```
- qos
- dot1p-classification classification-id [create]
- no dot1p-classification classification-id
  - description description-string
  - no description
  - dot1p dot1p-priority fc fc-name profile {in | out}
  - no dot1p dot1p-priority
- dscp-classification classification-id [use-ip-prec-classification] [create]
- no dscp-classification classification-id
  - description description-string
  - no description
  - dscp dscp-name fc fc-name profile {in | out}
  - no dscp dscp-name
  - prec ip-prec-value fc fc-name profile {in | out}
  - no prec ip-prec-value
- mpls-lsp-exp-classification classification-id [create]
- no mpls-lsp-exp-classification classification-id
  - description description-string
  - no description
  - lsp-exp mpls-lsp-exp-priority fc fc-name profile {in | out}
  - no lsp-exp mpls-lsp-exp-priority
```

6.3.1.3 Show commands

```
- show
- qos
- dot1p-classification policy-id association [detail]
- dscp-classification policy-id association [detail]
- mpls-lsp-exp-classification policy-id association [detail]
```

6.3.1.4 Operational commands

```
- config
- qos
- copy dot1p-classification src-pol dst-pol [overwrite]
- copy dscp-classification src-pol dst-pol [overwrite]
```

6.3.2 Command description

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

6.3.2.1 Configuration commands

dscp-classification

Syntax

dscp-classification *classification-id* [**use-ip-prec-classification**] [**create**]

no dscp-classification *classification-id*

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the mapping of the IP DSCP value in the IP header of the received packet to the forwarding class (FC) and ingress profile.



Note:

A DSCP classification policy with IP precedence can be configured only in a service egress policy and used only with access SAP egress.

The **no** form of this command removes the definition of the DSCP classification policy. The **no** form of the policy cannot be executed if the policy is in use; for example, if the policy is associated with a network QoS policy.

Default

dscp-classification 1

Parameters

classification-id

Specifies the DSCP classification policy.

Values 1 to 65535

use-ip-prec-classification

Keyword to configure the software to use IP precedence entries specified in the DSCP classification policy. If this keyword is not specified, the software uses DSCP classification entries. To maintain backward compatibility, **use-ip-prec-classification** is not configured by default.

This keyword is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

create

Keyword to create the DSCP classification policy.

description

Syntax

description *description-string*

no description

Context

```
config>qos>dscp-classification  
config>qos> dot1p-classification  
config>qos> mpls-lsp-exp-classification
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

dscp

Syntax

```
dscp dscp-name fc fc-name profile {in | out}  
no dscp dscp-name
```

Context

```
config>qos>dscp-classification
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command maps the IP DSCP value to an FC and assigns the ingress profile to the packet. The configured value is used to match the value in the received packet and assign the configured FC and profile on a exact match.

The **no** form of this command removes the mapping of the IP DSCP value to the FC.



Note:

A default FC is not assigned on executing the no form. The default FC is assigned in the QoS policy with which this policy is associated with.

Parameters

dscp dscp-name

Specifies the IP DSCP value to match by configuring the IP DSCP name corresponding to the DSCP value.

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 a case-sensitive, system-defined FC name.

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

profile {in | out}

Specifies whether the packets assigned to this FC are considered in or out of profile. A value of in defines the packet profile as in-profile and a value of out defines the packet profile to be out-of-profile.

prec

Syntax

prec *ip-prec-value* [**fc** *fc-name*] [**profile** {in | out}]

no prec

Context

config>qos>dscp-classification

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command defines a specific IP precedence value that must be matched to perform the associated reclassification actions. If an egress packet on the SAP matches the specified IP precedence value, the FC or profile may be overridden. By default, the FC and profile of the packet are derived from ingress classification and profiling functions.

The IP precedence bits used to match against precedence reclassification rules come from the Type of Service (ToS) field within the IPv4 header. If the packet does not have an IPv4 header, precedence-based matching is not performed.

The reclassification actions from an IP precedence reclassification rule may be overridden by a DSCP matching event.

The fc keyword is optional. When specified, the egress classification rule overwrites the FC derived from ingress. The new FC is used for egress remarking and queue mapping decisions. If a DSCP match occurs

after the IP precedence match, the new FC may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new FC, the FC from the IP precedence match is used.

The **profile** keyword is optional. When specified, the egress classification rule overwrites the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior. If a DSCP match occurs after the IP precedence match, the new profile may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new profile, the profile from the IP precedence match is used.

The **no** form of this command removes the reclassification rule from the SAP egress QoS policy.

Parameters

ip-prec-value

Specifies the IP precedence value.

Values 0, 1, 2, 3, 4, 5, 6, 7

fc-name

Specifies that packets matching the IP precedence value are explicitly reclassified to the specified FC regardless of the ingress classification decision. The explicit FC reclassification may be overwritten by a higher priority DSCP reclassification match. The FC name defined must be one of the eight FCs supported by the system. To remove the FC reclassification action for the specified precedence value, the **prec** command must be re-executed without the *fc-name* parameter defined.

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

profile

Optional keyword to specify that packets matching the IP precedence value are explicitly reclassified to the specified profile regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by a higher priority DSCP reclassification match. To remove the profile reclassification action for the specified precedence value, the **prec** command must be re-executed without the profile parameter defined.

in

Keyword to specify that any packets matching the reclassification rule are treated as in-profile by the egress forwarding plane.

out

Keyword to specify that any packets matching the reclassification rule are treated as out-of-profile by the egress forwarding plane.

dot1p-classification

Syntax

dot1p-classification *classification-id* [**create**]

no dot1p-classification *classification-id*

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the map of the dot1p value 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 policy. The **no** form of the policy cannot be executed if the policy is in use; for example, if the policy is associated with a network QoS policy.

Default

dot1p-classification 1

Parameters

classification-id

Specifies the policy ID.

Values 1 to 65535

create

Keyword to create.

dot1p

Syntax

dot1p *dot1p-priority* **fc** *fc-name* **profile** {**in** | **out**}

no dot1p *dot1p-priority*

Context

config>qos>dot1p-classification

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command maps the dot1p value to an FC and assigns the ingress profile to the packet. The configured value is used to match the value in the received packet and assign the configured FC and profile on a exact match.

The **no** form of this command removes the mapping of the dot1p value to FC.

A default FC is not assigned on executing the **no** form of this command. The default FC is assigned in the QoS policy with which this policy is associated with.

Parameters

fc fc-name

Specifies a case-sensitive, system-defined FC name.

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

profile {in | out}

Specifies whether the packets assigned to this FC are considered in or out of profile. A value of in defines the packet profile as in-profile and a value of out defines the packet profile to be out-of-profile.

dot1p-priority

Specifies the dot1p priority value to match.

Values 0 to 7

mpls-lsp-exp-classification

Syntax

mpls-lsp-exp-classification *classification-id* [create]

no mpls-lsp-exp-classification *classification-id*

Context

config>qos

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command defines the map of the MPLS EXP value in the MPLS header of the received packet to the FC and ingress profile.

The **no** form of this command removes the definition of the policy. The no form of the policy cannot be executed if the policy is in use; for example, if the policy is associated with a network QoS policy.

Default

default policy 1

Parameters

classification-id

Specifies the policy.

Values 1 to 65535

create

Keyword to create.

lsp-exp

Syntax

lsp-exp *mpls-lsp-exp-priority* **fc** *fc-name* **profile** {**in** | **out**}
no lsp-exp *mpls-lsp-exp-priority*

Context

config>qos>mpls-lsp-exp-classification

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command is used to map the MPLS EXP value to an FC and assigns the ingress profile to the packet. The configured value is used to match the value in the received packet and to assign the configured FC and profile on a exact match.

The **no** form of this command removes the mapping of the MPLS EXP value to the FC.



Note:

A default FC is not assigned on executing the **no** form. The default FC is assigned in the QoS policy with which this policy is associated with.

Default

default FC is not assigned; the mapping must be explicitly configured

Parameters

mpls-lsp-exp-priority

Specifies the LSP EXP value to match.

Values 0 to 7

fc fc-name

Specifies a case-sensitive, system-defined FC name.

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

profile {in | out}

Specifies whether the packets assigned to this FC is considered in or out of profile. A value of in defines the packet profile as in-profile and a value of out defines the packet profile to be out-of-profile.

6.3.2.2 Show commands

dot1p-classification

Syntax

```
dot1p-classification [policy-id] association
dot1p-classification [policy-id] detail
```

Context

```
show>qos
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the dot1p classification.

Parameters

- policy-id**
Displays information about the specific policy ID.
- associations**
Displays the associations of the dot1p classification policy.
- detail**
Displays the detailed information of the dot1p classification policy.

Output

The following output is an example of dot1p classification policy information, and [Table 29: Output fields: dot1p classification](#) describes the output fields.

Sample output

```
*A:Dut-A# show qos dot1p-classification 10 detail
=====
DOT1P Classification Maps
=====
-----
Dot1P Class Id      : 10
Description         : (Not Specified)
-----
Dot1P Bit Map              Forwarding Class              Profile
-----
No Matching Entries

-----
Network Policy Associations
-----
No Network Policy Associations found.
```

```
-----
SAP Ingress Associations
-----
No SAP Ingress Associations found.
-----
SAP Egress Associations
-----
SAP Egress Id           : 10
=====
```

Table 29: Output fields: dot1p classification

Label	Description
Dot1p Class Id	Displays the dot1p classification policy identifier
Description	Displays a text string that helps identify the policy context in the configuration file
Dscp Bit Map	Displays the dot1p value
Forwarding Class	Displays the forwarding class assigned
Profile	Displays the profile assigned
SAP Egress Id	Displays the associated SAP egress policy ID

dscp-classification

Syntax

```
dscp-classification [policy-id] association
dscp-classification [policy-id] detail
```

Context

```
show>qos
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the DSCP classification.

Parameters

```
policy-id
    Displays information about the specific policy ID.
```

associations

Displays the associations of the dscp-classification policy.

detail

Displays the detailed information of the dscp-classification policy.

Output

The following outputs are examples of DSCP classification policy information, and [Table 30: Output fields: DSCP classification](#) describes the output fields.

- [Sample output 1](#)
- [Sample output 2](#)

Sample output 1

```
*A:Dut-A# show qos dscp-classification 20 detail
=====
DSCP Classification Maps
=====
-----
Dscp Class Id      : 20
Description        : (Not Specified)
-----
Dscp Bit Map              Forwarding Class              Profile
-----
cp7                        nc                          None
-----
IP Prec Bit Map          Forwarding Class              Profile
-----
No Matching Entries
-----
Network Policy Associations
-----
No Network Policy Associations found.
-----
SAP Ingress Associations
-----
No SAP Ingress Associations found.
-----
SAP Egress Associations
-----
SAP Egress Id          : 10
=====
```

Sample output 2

```
*A:Dut-A# show qos dscp-classification 10 detail
=====
DSCP Classification Maps
=====
-----
Dscp Class Id      : 10
Description        : (Not Specified)
-----
Dscp Bit Map              Forwarding Class              Profile
-----
```

```

-----
No Matching Entries
-----
IP Prec Bit Map          Forwarding Class          Profile
-----
0                        nc                        In
1                        be                        None
2                        af                        Out
-----
Network Policy Associations
-----
No Network Policy Associations found.
-----
SAP Ingress Associations
-----
No SAP Ingress Associations found.
-----
SAP Egress Associations
-----
SAP Egress Id           : 10
=====

```

Table 30: Output fields: DSCP classification

Label	Description
Dscp Class Id	Displays the DSCP classification policy identifier
Description	Displays a text string that helps identify the policy context in the configuration file
Dscp Bit Map	Displays the DSCP value
Forwarding Class	Displays the forwarding class assigned
Profile	Displays the profile assigned
IP Prec Bit Map	Displays the IP precedence value This field is supported only on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
SAP Egress Id	Displays the associated SAP egress policy ID

mpls-lsp-exp-classification

Syntax

mpls-lsp-exp-classification [*policy-id*] **association**

mpls-lsp-exp-classification [*policy-id*] **detail**

Context

show>qos

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays the association between network ingress MPLS LSP-EXP values and FCes and profiles.

Parameters

- policy-id**
Displays information about the specific policy ID.
- associations**
Displays the associations of the mpls-lsp-exp-classification policy.
- detail**
Displays the detailed information of the mpls-lsp-exp-classification policy.

Output

The following output is an example of MPLS LSP-EXP information, and [Table 31: Output fields: MPLS LSP EXP classification](#) describes the output fields.

Sample output

```
=====
MPLS-LSP-EXP Classification Maps
=====
-----
Mpls-Lsp-Exp Class Id      : 10
Description                : (Not Specified)
-----
Mpls-Lsp-Exp Bit Map      Forwarding Class      Profile
-----
1                          nc                    In
-----
Network Policy Associations
-----
Network Policy Id          : 10
-----
SAP Ingress Associations
-----
No SAP Ingress Associations found.
=====
```

Table 31: Output fields: MPLS LSP EXP classification

Label	Description
Mpls-Lsp-Exp Class Id	Displays the MPLS LSP EXP classification policy identifier
Description	Displays a text string that helps identify the policy context in the configuration file
Mpls-Lsp-Exp Bit Map	Displays the MPLS LSP EXP value
Forwarding Class	Displays the forwarding class assigned
Profile	Displays the profile assigned
Network Policy Id	Displays the network policy identifier
SAP Ingress Associations	Displays the associated SAP ingress policy identifier

7 Network QoS policies

This chapter provides information to configure network QoS policies using the CLI.

7.1 Overview of network QoS policy on 7210 SAS-K 2F1C2T

Network QoS policies have an ingress and egress component, which define the QoS processing behavior to be provided for packets that ingress the access-uplink port and egress the access-uplink port respectively.

The ingress component of the policy defines how the IP DSCP and dot1p values using the DSCP and dot1p classification policies are mapped to internal FC and profile state for the 7210 SAS-K 2F1C2T. The FC 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 bandwidth-limiting parameters for the traffic mapped to each FC. Traffic mapped to each FC can be limited to configurable bandwidth values using separate queues for unicast traffic and multipoint traffic.

7210 SAS platforms provide different mechanisms to limit the bandwidth per FC. On 7210 SAS-K 2F1C2T, users can use the queue with packet buffers and a rate shaper to limit and shape the traffic per FC. Use of queue with shapers typically allows for better TCP traffic behavior in the network.

The egress component of the network QoS policy defines the marking values associated with each FC.

On the 7210 SAS, the user has an option to define the number of queues to use per access-uplink port and map the FC to queues. By default, network QoS policy 1 is used for access-uplink ports, until an explicit policy is associated. The default policy creates eight queues per access-uplink port. The queues are assigned default values for all the parameters defined with the default policy.

Access-uplink port egress marking supports the following:

- For packets sent out of an access-uplink port, the network QoS policy defines the marking values (for example, IEEE 802.1p bits and so on) to use based on the FC and the profile state.
- The default map of FC to marking values (for example, 802.1p bits and so on) is as shown in the default network QoS policy, *policy-id* 1.
- All non-default network QoS policies inherit the default map and can be modified by the user.
- Remarking can be enabled or disabled on access-uplink ports.
- An option is available to map FC and profile to either IP DSCP and dot1p bits along with DEI bit.

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 access uplink ports 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, and then the original policy-id can be overwritten using the **config qos copy** command.

See the CLI usage chapter in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for information about the tasks and commands required to access the CLI and to configure and maintain 7210 SAS devices.

7.1.1 Resource allocation for network QoS policy for 7210 SAS-K 2F1C2T

The queues required for access-uplink port egress are allocated for the egress queue system pool. Queues from the egress queue system pool are allocated for per SAP egress queues and per access-uplink port egress queues.

The dot1p policy and DSCP policy resources used for network qos ingress FC assignment on access-uplink port, is shared with the per SAP ingress classification criteria from the system pool.

7.2 Basic configurations for 7210 SAS-K 2F1C2T

A basic network QoS policy must conform to the following:

- have a unique policy ID
- specify the default-action
- have a QoS policy scope of **template** or **exclusive**
- on a 7210 SAS platform, have at least one default unicast FC meter/queue
- on a 7210 SAS platform, have at least one multipoint FC meter/queue

7.2.1 Create a network QoS policy for access-uplink ports on 7210 SAS-K 2F1C2T

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

To create an network QoS policy, define the following:

- Specify a network policy ID value. The system does not dynamically assign a value.
- Include a description that provides a brief overview of policy features.
- Use egress marking and remarking to specify the egress FC to marking value (for example, IEEE 802.1p and so on) map; otherwise, the default values are applied. The following are defined:
 - **remarking**
If enabled, this command remarks ALL packets that egress on the specified access uplink port. The remarking is based on the FC to marking values mapping defined under the egress node of the network QoS policy. On the 7210 SAS-K 2F1C2T, the user has an option to enable it or disable it.
 - **FC criteria**
The FC name represents an egress queue. Specify FC criteria to define the marking criteria of packets flowing through it.
 - **marking value**
The marking (for example, IEEE 802.1p) value is used for all packets requiring marking that egress on this FC queue that are in or out of profile.

- Specify ingress criteria to use for FC mapping for all packets using the following:
 - **default action**
Defines the default action to be taken for packets that have undefined dot1p bits set. The default-action specifies the FC to which the packets are assigned.
 - **dot1p**
On the 7210 SAS-K 2F1C2T, the user has an option to specify either dot1p or IP DSCP to FC mapping for all packets. Ingress traffic that matches the specified criteria are assigned to the corresponding FC.

Use the following syntax to create a network QoS policy for 7210 SAS-K 2F1C2T.

```
config>qos#
config>qos>network network-policy id create
description description-string
scope {exclusive | template}
egress
  [no] remarking
  remark remark-policy id
  no remark
ingress
  default-action fc fc name
  dot1p-classification dot1p-classification id
  dscp-classification dscp-classification id
  queue queue id create
  fc fc-name create
    queue queue-id
    multicast-queue queue-id
    [no] use-dei
```

Use the following syntax to associate a network QoS policy with the access-uplink port.

```
config>port
  ethernet
    access
      uplink
        qos network-policy-id
```

```
config>router
  interface interface-name
    qos network-policy-id
```

Example

The following is a sample configuration output that displays uplink port 1/1/1 with network policy 600 applied to the interface.

```
A:ALA-7>config# info
#-----
echo "Port Configuration"
#-----
  port 1/1/1
    shutdown
    description "port 1/1/1"
    ethernet
      mode access uplink
      access
        uplink
```

```

                                qos 600
                                exit
                            exit
                        exit
                    exit
                ...
            #-----
A:ALA-7>config#

```

7.2.1.1 Default network policy values for access-uplink ports on 7210 SAS-K 2F1C2T

The default network policy for access-uplink ports is identified as policy-id "1". Default policies cannot be modified or deleted.

Example

The following is sample configuration output that shows the default network policy parameters.

```

7210SAS>config>qos>network# info
-----
description "Default network-port QoS policy."
ingress
    dot1p-classification 1
    queue 1 create
    exit
    queue 2 create
        rate cir 25
    exit
    queue 3 create
        rate cir 25
    exit
    queue 4 create
        rate cir 25
    exit
    queue 5 create
        rate cir 100
    exit
    queue 6 create
        rate cir 100
    exit
    queue 7 create
        rate cir 10
    exit
    queue 8 create
        rate cir 10
    exit
    fc "af" create
        queue 3
        multicast-queue 3
    exit
    fc "be" create
        queue 1
        multicast-queue 1
    exit
    fc "ef" create
        queue 6
        multicast-queue 6
    exit
    fc "h1" create
        queue 7
        multicast-queue 7

```

```

        exit
        fc "h2" create
            queue 5
            multicast-queue 5
        exit
        fc "l1" create
            queue 4
            multicast-queue 4
        exit
        fc "l2" create
            queue 2
            multicast-queue 2
        exit
        fc "nc" create
            queue 8
            multicast-queue 8
        exit
    exit
    egress
    exit
-----
*7210 SAS>config>qos>network#

```

The following table list default network policy parameters.

Table 32: Network policy ID #1 defaults for access-uplink ports

Field	Default
description	Default network QoS policy.
scope	template
ingress	
default-action	fc be profile out
egress	
remarking	No
fc af:	
dot1p-in-profile	3
dot1p-out-profile	2
fc be:	
dot1p-in-profile	0
dot1p-out-profile	0
fc ef:	
dot1p-in-profile	5
dot1p-out-profile	5

Field	Default
fc h1:	
dot1p-in-profile	6
dot1p-out-profile	6
fc h2:	
dot1p-in-profile	4
dot1p-out-profile	4
fc l1:	
dot1p-in-profile	3
dot1p-out-profile	2
fc l2:	
dot1p-in-profile	1
dot1p-out-profile	1
fc nc:	
dot1p-in-profile	7
dot1p-out-profile	7

Table 33: Default network QoS policy ID #1 for dot1p to FC mapping for 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

7.3 Overview of network QoS policy on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The network QoS policy has an ingress and egress component, which define the QoS processing behavior for packets that ingress into the network port, hybrid port, and access-uplink port, and egress from the network port, hybrid port, and access-uplink port, respectively.

7.3.1 Network QoS policy for access-uplink ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The ingress component of the policy defines how the IP DSCP and dot1p values using the DSCP and dot1p classification policies are mapped to the internal FC and profile state for the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C. The FC 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 bandwidth-limiting parameters for the traffic mapped to each FC. Traffic mapped to each FC can be limited to configurable bandwidth values using separate queues for unicast traffic and multipoint traffic.



Note:

The 7210 SAS platforms provide different mechanisms to limit the bandwidth per FC. On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, use the queue with packet buffers and a rate shaper to limit and shape the traffic per FC. Use of queues with shapers allows better TCP traffic behavior in the network.

The egress component of the network QoS policy defines the marking values associated with each FC.

The user has an option to define the number of queues to use per access-uplink port and map the FC to queues. By default, network QoS policy "1" is used for access-uplink ports, until an explicit policy is associated. The default policy creates eight queues per access-uplink port. The queues are assigned default values for all the parameters defined with the default policy.

Access-uplink port egress marking supports the following:

- For packets sent out of an access-uplink port, the network QoS policy defines the marking values (for example, IEEE 802.1p bits and so on) to use based on the FC and the profile state.
- The default map of FC to marking values (for example, 802.1p bits) is as shown in the default network QoS policy, *policy-id* 1.
- All non-default network QoS policies inherit the default map and can be modified by the user.
- Remarking can be enabled or disabled on access-uplink ports.
- An option is available to map FC and profile to either IP DSCP and dot1p bits along with DEI bit.

Non-default network policy parameters can be modified. The **no** form of the command reverts to the default values.

Changes made to a policy are applied immediately to all access uplink ports 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, and then the original *policy-id* can be overwritten using the **config qos copy** command.

See the CLI usage chapter in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for information about the tasks and commands required to access the CLI, and to configure and maintain your devices.

7.3.2 Network QoS policy for network ports and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The ingress component of the policy defines how the MPLS EXP, IP DSCP, and dot1p values that use classification policies are mapped to internal FC and profile state for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C. The FC and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the system. The mapping on each network port and hybrid port defaults to the mapping defined in the default network QoS policy, unless an explicit policy is defined for the network ports. The policy 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 queues for unicast traffic and multipoint traffic.



Note:

The 7210 SAS platforms provide different mechanisms to limit the bandwidth per FC. On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, use the queue with packet buffers and a rate shaper to limit and shape the traffic per FC. Use of queues with shapers allows better TCP traffic behavior in the network.

The egress component of the network QoS policy defines the marking values associated with each FC.

The user has an option to define the number of queues to use per network port and hybrid port, and to map the FC to queues. By default, network QoS policy "2" is used for network ports, unless an explicit policy is associated. The default policy creates eight queues per network port or hybrid port. The queues are assigned default values for all the parameters defined with the default policy.

Network port egress marking supports the following:

- For packets sent out of the IP interfaces that are configured on a network port or hybrid port, the network QoS policy defines the marking values (for example, IEEE 802.1p bits, and so on) to use based on the FC and the profile state.
- The default map of FC to marking values (for example, 802.1p bits) is as shown in default network QoS policy, *policy-id 2* for network port.
- All non-default network QoS policies inherit the default map and can be modified by the user.
- Remarking can be enabled or disabled on access-uplink ports.
- An option is available to map FC and profile as follows:
 - for MPLS packets - option to mark MPLS EXP values and dot1p bits along with DEI bit
 - for IP packets - option to mark IP DSCP and dot1p bits along with DEI bit

Non-default network policy parameters can be modified. The **no** form of the command reverts to the default values.

Changes made to a policy are applied immediately to all access-uplink ports 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, and then the original *policy-id* can be overwritten using the **config qos copy** command.

See the CLI usage chapter in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for information about the tasks and commands required to access the CLI, and to configure and maintain your devices.

7.3.3 Resource allocation for network QoS policy for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The queues required for network port egress, hybrid port egress, and access-uplink port egress are allocated from the egress queue system pool. Queues from the egress queue system pool are shared among SAP egress queues, access-uplink port egress queues, network port egress queues, and hybrid port egress queues.

The dot1p policy and DSCP policy resources used for network QoS ingress FC assignment on network ports, hybrid ports, and access-uplink ports are shared with the per-SAP ingress classification criteria from the system pool.

The MPLS EXP policy resources used for network QoS ingress FC assignment on network ports and hybrid ports are allocated from the system pool.

7.4 Basic configurations for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

A basic network QoS policy must conform to the following:

- have a unique policy ID
- specify the default-action
- have a QoS policy scope of **template** or **exclusive**
- on a 7210 SAS platform, have at least one default unicast FC queue
- on a 7210 SAS platform, have at least one multipoint FC queue

7.4.1 Create a network QoS policy on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

This section provides information about creating a network QoS policy on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

7.4.1.1 Create a network QoS policy for access-uplink ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

See [Create a network QoS policy for access-uplink ports on 7210 SAS-K 2F1C2T](#) for more information.

7.4.1.1.1 Default network policy values for access-uplink ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The default network policy for access-uplink ports is identified as *policy-id* 1. Default policies cannot be modified or deleted.

Example

The following is a sample configuration output for default network policy ID 1 used for the access-uplink port on the 7210 SAS-K 2F6C4T.

```
*A:K-SASK12>config>qos>network# info
-----
description "Default network-port QoS policy."
ingress
  dot1p-classification 1
  queue 1 create
  exit
  queue 2 create
    rate cir 25
  exit
  queue 3 create
    rate cir 25
  exit
  queue 4 create
    rate cir 25
  exit
  queue 5 create
    rate cir 100
  exit
  queue 6 create
    rate cir 100
  exit
  queue 7 create
    rate cir 10
  exit
  queue 8 create
    rate cir 10
  exit
  fc "af" create
    queue 3
    multicast-queue 3
  exit
  fc "be" create
    queue 1
    multicast-queue 1
  exit
  fc "ef" create
    queue 6
    multicast-queue 6
  exit
  fc "h1" create
    queue 7
    multicast-queue 7
  exit
  fc "h2" create
    queue 5
    multicast-queue 5
  exit
  fc "l1" create
    queue 4
    multicast-queue 4
  exit
  fc "l2" create
    queue 2
    multicast-queue 2
  exit
  fc "nc" create
    queue 8
```



```

        multicast-queue 8
    exit
    exit
    egress
    exit
-----
*A:K-SASK12>config>qos>network#

```

Example

The following is a sample configuration output for default network policy ID 1 used for the access-uplink port on the 7210 SAS-K 3SFP+ 8C.

```

*A:Dut-A>config>qos>network# info detail
-----
description "Default network-port QoS policy."
scope template
ingress
    default-action fc be
    dot1p-classification 1
    no dscp-classification
    no mpls-lsp-exp-classification
    queue 1 create
        rate cir 0 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    queue 2 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    queue 3 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    queue 4 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    queue 5 create
        rate cir 15 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24

```

```
        weight 1
        priority 1
        slope-policy "default"
    exit
queue 6 create
    rate cir 15 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
queue 7 create
    rate cir 5 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
queue 8 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
fc "af" create
    queue 3
    multicast-queue 3
    no use-dei
exit
fc "be" create
    queue 1
    multicast-queue 1
    no use-dei
exit
fc "ef" create
    queue 6
    multicast-queue 6
    no use-dei
exit
fc "h1" create
    queue 7
    multicast-queue 7
    no use-dei
exit
fc "h2" create
    queue 5
    multicast-queue 5
    no use-dei
exit
fc "l1" create
    queue 4
    multicast-queue 4
    no use-dei
exit
fc "l2" create
    queue 2
    multicast-queue 2
```

```

        no use-dei
    exit
    fc "nc" create
        queue 8
        multicast-queue 8
        no use-dei
    exit
exit
egress
    no remarking
    remark 1
exit
-----

```

7.4.1.2 Create a network QoS policy for network ports and hybrid ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Configuring and applying QoS policies other than the default policy is optional. A default network policy is applied to each network port or hybrid port.

To create a network QoS policy, define the following:

- Specify a network policy ID value. The system does not dynamically assign a value.
- Include a description that provides a brief overview of policy features.
- Use egress marking and remarking to specify the egress FC to marking value (for example, IEEE 802.1p and so on) map; otherwise, the default values are applied. The following are defined:
 - **marking**
If enabled, the system remarks packets that egress the IP interfaces configured on the network port or hybrid port. The remarking is based on the FC to marking values mapping defined under the egress node of the network QoS policy. The user has an option to enable or disable remarking.
 - **FC criteria**
The FC name represents an egress queue. Specify FC criteria to define the marking criteria of packets flowing through it.
 - **marking value**
The marking (for example, IEEE 802.1p) value is used for all packets requiring marking that egress on this FC queue that are in or out of profile.
- Specify the ingress criteria to use for FC mapping for all packets using the following:
 - **default action**
Defines the default action to be taken for packets that have undefined dot1p bits set. The default-action specifies the FC to which the packets are assigned.
 - The user has an option to specify either MPLS EXP and dot1p and IP DSCP to FC mapping for all packets. Ingress traffic that matches the specified criteria are assigned to the corresponding FC.

Use the following syntax to create a network QoS policy for network port on the 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C.

config>qos>network

network *network-policy-id* [**create**]

no network *network-policy-id*

Example

The following commands associate a network QoS policy with the network port on the 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C.

```
*A:K-SASK12>config>qos>network# info
-----
description "Default network QoS policy."
ingress
  dot1p-classification 1
  dscp-classification 1
  mpls-lsp-exp-classification 1
  queue 1 create
  exit
  queue 2 create
    rate cir 25
  exit
  queue 3 create
    rate cir 25
  exit
  queue 4 create
    rate cir 25
  exit
  queue 5 create
    rate cir 100
  exit
  queue 6 create
    rate cir 100
  exit
  queue 7 create
    rate cir 10
  exit
  queue 8 create
    rate cir 10
  exit
  fc "af" create
    queue 3
    multicast-queue 3
  exit
  fc "be" create
    queue 1
    multicast-queue 1
  exit
  fc "ef" create
    queue 6
    multicast-queue 6
  exit
  fc "h1" create
    queue 7
    multicast-queue 7
  exit
  fc "h2" create
    queue 5
    multicast-queue 5
  exit
  fc "l1" create
    queue 4
    multicast-queue 4
  exit
  fc "l2" create
    queue 2
    multicast-queue 2
  exit
```

```

        fc "nc" create
        queue 8
        multicast-queue 8
    exit
exit
egress
    remark 2
exit
-----
*A:K-SASK12>config>qos>network#

```

Example

The following is a sample configuration output for network port 1/1/1 with network policy 600 applied to the port.

```

#-----
echo "Port Configuration"
#-----
    port 1/1/1
        shutdown
        ethernet
            network
                qos 2
            exit
        exit
    exit
exit
    port 1/1/2
        shutdown
        ethernet
            network
                qos 2
            exit
        exit
    exit
exit
    port 1/1/3
        shutdown
        ethernet
            network
                qos 2
            exit
        connection-type copper
    exit
exit
    port 1/1/4
        shutdown
        ethernet
            network
                qos 2
            exit
        exit
    exit
exit
    port 1/1/5
        ethernet
            mode access
            access
            exit
            mtu 1518
        exit
        no shutdown
    exit
#-----

```

7.4.1.3 Default network policy values on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The default network policy for network ports is identified as policy-id "2". Default policies cannot be modified or deleted.

Example

The following is a sample configuration output for default network policy ID 2 used for the network port on the 7210 SAS-K 2F6C4T.

```
*A:K-SASK12>config>qos>network# info detail
-----
description "Default network QoS policy."
scope template
ingress
  default-action fc be
  dot1p-classification 1
  dscp-classification 1
  mpls-lsp-exp-classification 1
  queue 1 create
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
  exit
  queue 2 create
    rate cir 25 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
  exit
  queue 3 create
    rate cir 25 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
  exit
  queue 4 create
    rate cir 25 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
  exit
  queue 5 create
    rate cir 100 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
```

```
        slope-policy "default"
    exit
    queue 6 create
        rate cir 100 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    queue 7 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    queue 8 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    fc "af" create
        queue 3
        multicast-queue 3
        no use-dei
    exit
    fc "be" create
        queue 1
        multicast-queue 1
        no use-dei
    exit
    fc "ef" create
        queue 6
        multicast-queue 6
        no use-dei
    exit
    fc "h1" create
        queue 7
        multicast-queue 7
        no use-dei
    exit
    fc "h2" create
        queue 5
        multicast-queue 5
        no use-dei
    exit
    fc "l1" create
        queue 4
        multicast-queue 4
        no use-dei
    exit
    fc "l2" create
        queue 2
        multicast-queue 2
        no use-dei
    exit
```

```

        fc "nc" create
            queue 8
            multicast-queue 8
            no use-dei
        exit
    exit
    egress
        no remarking
        remark 2
    exit
-----
*A:K-SASK12>config>qos>network#

```

Example

The following is a sample configuration output for default network policy ID 2 used for the network port on the 7210 SAS-K 3SFP+ 8C.

```

*A:Dut-A>config>qos>network# info detail
-----
description "Default network QoS policy."
scope template
ingress
    default-action fc be
    dot1p-classification 1
    dscp-classification 1
    mpls-lsp-exp-classification 1
    queue 1 create
        rate cir 0 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    queue 2 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    queue 3 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
    queue 4 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit

```



```
queue 5 create
    rate cir 15 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
queue 6 create
    rate cir 15 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
queue 7 create
    rate cir 5 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
queue 8 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
fc "af" create
    queue 3
    multicast-queue 3
    no use-dei
exit
fc "be" create
    queue 1
    multicast-queue 1
    no use-dei
exit
fc "ef" create
    queue 6
    multicast-queue 6
    no use-dei
exit
fc "h1" create
    queue 7
    multicast-queue 7
    no use-dei
exit
fc "h2" create
    queue 5
    multicast-queue 5
    no use-dei
exit
fc "l1" create
    queue 4
    multicast-queue 4
```

```

        no use-dei
    exit
    fc "l2" create
        queue 2
        multicast-queue 2
        no use-dei
    exit
    fc "nc" create
        queue 8
        multicast-queue 8
        no use-dei
    exit
exit
egress
    no remarking
    remark 2
exit
-----

```

7.5 DSCP marking CPU self-generated traffic

The user can configure DSCP marking for CPU-generated traffic on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms only. See [QoS for CPU self-generated traffic on network interfaces for the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#) for more information about QoS for self-generated (CPU) traffic on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C. The following table lists the default DSCP and dot1p marking values for the 7210 SAS-K 2F1C2T.



Note:

Protocols such as BGP, RSVP, TLDP, OSPF, and IS-IS are not supported on 7210 SAS devices configured in access-uplink mode.

Table 34: DSCP and dot1p marking for 7210 SAS-K 2F1C2T

Protocol	IPv4	DSCP marking	dot1p marking	Default FC	DSCP values	dot1p values
ARP	NA	NA	Yes	NC	-	7
CFM	NA	NA	Yes	NC	-	7
FTP	Yes	Yes	Yes	H2	34	4
ICMP ping	Yes	Yes	Yes	NC	0	7
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
NTP	Yes	Yes	Yes	H2	48	7

Protocol	IPv4	DSCP marking	dot1p marking	Default FC	DSCP values	dot1p values
PTP	Yes	Yes	Yes	H2	48	7
RADIUS	Yes	Yes	Yes	H2	34	4
SCP	NA	NA	Yes	H2	34	4
SNMP	Yes	Yes	Yes	H2	34	4
SNMP trap/ log	Yes	Yes	Yes	H2	34	4
SSH	Yes	Yes	Yes	H2	34	7
STP	NA	NA	Yes	NC	-	7
SYSLOG	Yes	Yes	Yes	H2	34	4
TACACS	Yes	Yes	Yes	H2	34	4
TACPLUS	Yes	Yes	Yes	H2	34	4
TELNET	Yes	Yes	Yes	H2	34	4
TFTP	Yes	Yes	Yes	H2	34	4
Trace route	Yes	Yes	Yes	NC	0	7

The following table lists the default DSCP and dot1p marking values for the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.



Note:

DSCP and dot1p values in [Table 35: DSCP and dot1p marking for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#) are applicable when remarking is disabled on the egress service object (for example, access port, SAP, and network port) used to send out the packets.

Table 35: DSCP and dot1p marking for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Protocol	IPv4	DSCP marking	dot1p marking	Default FC	DSCP values (decimal)	dot1p values
OSPF	Yes	Yes	Yes	NC	48	7
ISIS	Yes	No	Yes	NC	-	7
TLDP	Yes	Yes	Yes	NC	48	7
RSVP	Yes	Yes	Yes	NC	48	7
SNMP	Yes	Yes	Yes	H2	34	4

Protocol	IPv4	DSCP marking	dot1p marking	Default FC	DSCP values (decimal)	dot1p values
NTP	Yes	Yes	Yes	NC	48	7
TELNET	Yes	Yes	Yes	H2	34	4
FTP	Yes	Yes	Yes	H2	34	4
TFTP	Yes	Yes	Yes	H2	34	4
SYSLOG	Yes	Yes	Yes	H2	34	4
TACACS	Yes	Yes	Yes	H2	34	4
RADIUS	Yes	Yes	Yes	NC	48	7
SSH	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
SCP	Yes	Yes	Yes	H2	34	4
PIM (SSM)	Yes	Yes	Yes	NC	48	7
STP	NA	NA	Yes	NC	-	7
CFM	NA	NA	Yes	NC	-	7
ARP	NA	NA	Yes	NC	-	7
Trace route	Yes	Yes	Yes	NC	0	7
TACPLUS	Yes	Yes	Yes	H2	34	4
IGMP	Yes	Yes	Yes	NC	48	7
DNS	Yes	Yes	Yes	H2	34	4
BGP	Yes	Yes	Yes	NC	48	7
PTP (see note) ⁹	Yes	Yes	Yes	see note ⁹	see note ⁹	7

⁹ Based on the type of the PTP message, that is, PTP event messages (for example, Sync message) and PTP non-event messages (for example, Announce, Follow-up), the DSCP value used is either 0x30 (h1) or 0x38 (nc), and the dot1p value is always 7.

7.5.1 QoS for CPU self-generated traffic on network interfaces for the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

DSCP, 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.

DSCP marking for internally-generated control and management traffic should be used for a specified application. The DSCP marking 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 also be set per application. When an application is configured to use a specified DSCP value and 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, because it applies to the logical interface the packet is egressing.

Configuring self-generated QoS traffic is supported in the base router and VPRN service contexts.

[Table 34: DSCP and dot1p marking for 7210 SAS-K 2F1C2T](#) and [Table 35: DSCP and dot1p marking for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#) list the default values for self-generated traffic on network interfaces.



Note:

- ICMP echo requests (type 8) and ICMPv6 echo requests (type 128) initiated from the router use the DSCP value set using the **ping** command TOS value.
- 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.

7.5.2 Default DSCP mapping

The following table lists the DSCP mapping between DSCP names and DSCP values (decimal, hexadecimal, and binary), and labels.

Table 36: 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

DSCP name	DSCP value decimal	DSCP value hexadecimal	DSCP value binary	Label
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	0b0100010	h2
af42	36	0x24	0b100100	h2
af43	38	0x26	0b100110	h2
default ¹⁰	0			

7.6 Service management tasks

This section provides information about service management tasks.

7.6.1 Deleting QoS policies

A network policy is associated by default with access-uplink ports.

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.

7.6.2 Remove a policy from the QoS configuration

Use the following syntax to delete a network policy.

```
config>qos# no network network-policy-id
```

¹⁰ The default FC mapping is used for all DSCP names/values for which there is no explicit FC mapping.

7.6.3 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.

Use the following syntax to overwrite a network policy.

```
config>qos# copy network source-policy-id dest-policy-id [overwrite]
```

Example: 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#
```

7.6.4 Editing QoS policies

You can change existing policies, except the default policies, and entries in the CLI. The changes are applied immediately to all 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.

7.7 Network QoS policy command reference

- [Command hierarchies](#)
- [Command descriptions](#)

7.7.1 Command hierarchies

- [Configuration commands for 7210 SAS-K 2F1C2T](#)

- [Configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [Self-generated traffic commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [Operational commands](#)
- [Show commands](#)

7.7.1.1 Configuration commands for 7210 SAS-K 2F1C2T

```
- config
- qos
- [no] network network-policy-id [create]
- description description-string
- no description
- scope {exclusive | template}
- no scope
- egress
- remark policy-id
- no remark
- [no] remarking
- [no] fc fc-name
- ingress
- default-action fc fc-name profile {in | out | use-dei}
- dot1p-classification policy-id
- no dot1p-classification
- dscp-classification policy-id
- no dscp-classification
- [no] fc fc-name
- use-dei
- no use-dei
- queue queue-id [create]
- no queue queue-id
- [no] adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
- cbs size-in-kbyte
- no cbs
- mbs size-in-kbytes
- no mbs
- priority level
- no priority
- rate [cir cir-percent] [pir pir-percent]
- no rate
- slope-policy name
- no slope-policy
- weight weight
- no weight
```

7.7.1.2 Configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
- config
- qos
- [no] network network-policy-id [create]
- egress
- remark policy-id
- no remark
- [no] remarking
- [no] fc fc-name
- ingress
- default-action fc fc-name profile {in | out | use-dei}
```



```

- dot1p-classification policy-id
- no dot1p-classification
- dscp-classification policy-id
- no dscp-classification
- mpls-lsp-exp-classification policy-id
- no mpls-lsp-exp-classification
- [no] fc fc-name
  - use-dei
  - no use-dei
- queue queue-id [create]
- no queue queue-id
  - [no] adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
  - cbs size-in-kbyte
  - no cbs
  - mbs size-in-kbytes
  - no mbs
  - priority level
  - no priority
  - rate [cir cir-percent] [pir pir-percent]
  - no rate
  - slope-policy name
  - no slope-policy
  - weight weight
  - no weight
- scope {exclusive | template}
- no scope

```

7.7.1.3 Self-generated traffic commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```

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
      - 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

```

7.7.1.4 Operational commands

```

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

```

7.7.1.5 Show commands

```
- show
  - qos
    - network [network-policy-id] association
    - network [network-policy-id] [detail]
  - router [router-instance]
  - router service-name service-name
    - sgt-qos
      - application [app-name] [dscp | dot1p]
      - dscp-map [dscp-name]
```

7.7.2 Command descriptions

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

7.7.2.1 Configuration commands

- [Generic commands](#)
- [Operational commands](#)
- [Network QoS Policy commands](#)
- [Network egress QoS policy commands](#)
- [Network ingress QoS policy commands](#)
- [Network ingress queue QoS policy commands](#)
- [Self-generated traffic commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)

7.7.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>network

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

7.7.2.1.2 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

Description

This command copies existing QoS policy entries for a QoS policy ID to another QoS policy ID.

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

Parameters

network src-pol dst-pol

Specifies that the source and destination policies are network policy IDs. Specify the source policy that the **copy** command will copy and specify the destination policy to which the command will duplicate the policy to a new or different policy ID.

Values 1 to 65535

overwrite

Keyword to replace the existing destination policy. Everything in the existing destination policy is overwritten with the contents of the source policy. If **overwrite** is not specified, an error occurs if the destination policy ID exists.

remark

Syntax

remark *policy-id*

no remark

Context

config>qos>network>egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the remarking policy ID to use for marking packets on access-uplink port egress, network port egress, or hybrid port egress. The usage is as follows for different platforms.

- On the 7210 SAS-K 2F1C2T, this policy is used to configure marking for packets sent out of access-uplink ports.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, this policy is used to configure marking for packets sent out of access-uplink ports, network ports, or hybrid ports, depending on the port that the policy is associated with.

For access-uplink ports, the remarking policy ID must be associated with the network QoS policy that is associated with the access-uplink port. Remarking must be enabled in the network QoS policy to enable marking of packets sent out of an access-uplink port. Only a remarking policy of type dot1p, dscp, or dot1p-dscp is allowed to be used when the remark policy is associated with access-uplink port egress. See [Remark policies](#) for more information about remark policy types and its usage.

For network ports and hybrid ports, the remarking policy ID must be associated with the network QoS policy that is associated with the network port and hybrid port. Remarking must be enabled in the network QoS policy to enable marking of packets sent out of network IP interfaces that are configured on the network port or hybrid port. On network ports or hybrid ports, the dot1p bits are marked by default, irrespective of whether remarking is enabled or disabled. Only a remarking policy of type lsp-exp, dot1p, dscp, dot1p-dscp, dot1p-lsp-exp, or dot1p-dscp-lsp-exp is allowed when the remark policy is associated with network port egress. See [Remark policies](#) for more information about remark policy types and their usage.

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, then the default remark policy is used to mark sent packets. If no remark policy is executed and remarking is disabled, then packets are not remarked at all.

Parameters

policy-id

Specifies the remark policy.

Values 1 to 65535

remarking

Syntax

[no] remarking

Context

config>qos>network>egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configured marking for packets. The marking functionality is as follows:

- On the 7210 SAS-K 2F1C2T, this command enables marking for packets sent out of access-uplink ports.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, this command enables marking for packets sent out of access-uplink ports, network ports, or hybrid ports, depending on the port that the policy is associated with.

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 and remark policy of type dot1p is configured in the network QoS policy, then the FC to dot1p mapping is used to mark packets sent out of the port.

See [Remark policies](#) for more information about how to configure remark policies.

The **no** form of this command disables remarking.

Default

no remarking

scope

Syntax

scope {exclusive | template}

no scope

Context

config>qos>network

Platforms

Supported on all 7210 SAS platforms as described in this document

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 becomes available for assignment to another exclusive interface. The system default policies cannot be put into the exclusive scope. An error is 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. An error is generated if you try to modify the template scope parameter to exclusive scope on default policies.

7.7.2.1.3 Network QoS Policy commands

network

Syntax

network *network-policy-id* [**create**]

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 functionality is as follows:

- On the 7210 SAS-K 2F1C2T, the network policy defines the treatment the packets receive on ingress into and egress from the access-uplink port.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the network policy defines the treatment the packets receive on ingress into and egress from the access-uplink port, network port, and hybrid port.

On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, the network QoS policy associated with access-uplink port has ingress and egress components. These components are described as follows.

- The ingress component of the policy defines how dot1p bits or IP DSCP 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 the same or separate queue for unicast and multipoint traffic.
- The egress component of the network qos policy defines the forwarding class and profile to packet header priority bits; for example, dot1p bits. Option is provided to map forwarding class to dot1p bits and IP DSCP bits.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the network QoS policy associated with network ports or hybrid ports has ingress and egress components. These components are described as follows:

- The ingress component of the policy defines how MPLS EXP bits or dot1p bits or IP DSCP 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 the same or separate queue for unicast and multipoint traffic.
- The egress component of the network QoS policy defines the forwarding class and profile to packet header priority bits (for example, dot1p bits). The option is provided to map the forwarding class to MPLS EXP bits, dot1p bits and IP DSCP bits.

The default network policy ID 1 is associated with access-uplink ports that do not have an explicit user configured policy and cannot be modified or deleted. The default network policy ID 2 is associated with network ports and hybrid ports that do not have an explicit user configured policy and cannot be modified or deleted. The default network policies define default mapping for packet header bits to the FCs on ingress and the mapping of the FC to queues.

If a new network policy is created, only the default action, default queues 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 QoS bit to FC mapping (for example, dot1p-to-FC mapping or EXP to FC mapping) for network QoS policy. The default network policy can be copied using the **copy** command to create a new network policy that includes the default ingress dot1p 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, are applied immediately to all the 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 policy-id** 1 cannot be deleted.

Default

System Default Network Policy 1

Parameters

network-policy-id

Specifies the policy on the 7210 SAS.

Values 1 to 65535

7.7.2.1.4 Network egress QoS policy commands

egress

Syntax

egress

Context

config>qos>network

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context create or edit egress policy entries that specify the forwarding class to marking values maps to be instantiated when this policy is applied to the access-uplink port, network port, or hybrid port.

The forwarding class and profile states are mapped to marking values for all packets affected by the policy defined in this context.

fc

Syntax

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

Context

config>qos>network>egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the forwarding class name. The forwarding class name represents an egress queue. The 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.

The **no** form of this command removes the forwarding class to marking value association. The forwarding class reverts to the mapping defined in the default network policy.

Default

undefined forwarding classes default to the configured parameters in the default network policy ID 1

Parameters

fc-name

Specifies a case-sensitive, system-defined forwarding class name.

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

7.7.2.1.5 Network ingress QoS policy commands

ingress

Syntax

ingress

Context

config>qos>network

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context creates or edits policy entries that specify the QoS bits to forwarding class mapping for all packets.

When pre-marked packets ingress on a network port or hybrid port, the QoS treatment through the 7210 SAS is based on the mapping defined on 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

Description

This command defines or edits the default action to be taken for packets do not match any of the configured classification entries. The **default-action** command specifies the forwarding class to which such packets are assigned.

Multiple **default-action** commands 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 dot1p or dot1p bits that is not defined are 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.

Values in — Any packets matching the classification rule are treated as in-profile packets.
out — Any packets matching the classification rule are treated as out-of-profile packets.
use-dei — The DEI bit received in the Ethernet VLAN tag is used to determine the profile of the packets. The packet will be considered to be in-profile if the DEI bit value is 0 and out-of-profile if the DEI value is 1.

dot1p-classification

Syntax

dot1p-classification *policy-id*

no dot1p-classification

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command associates a dot1p classification policy to a forwarding class and profile state based on the dot1p bits in the packet. The dot1p classification policy contains entries used to map traffic received on

access-uplink ports, or traffic received in the context of a network IP interface configured on a network port or hybrid port.

The **no** form of this command disables the use of this policy.

Default

no dot1p-classification-policy

Parameters

policy-id

Specifies the policy on the 7210 SAS.

Values 1 to 65535

dscp-classification

Syntax

dscp-classification *policy-id*

no dscp-classification

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command associates an IP DSCP classification policy to a forwarding class and profile state based on the IP DSCP bits in the packet. The IP DSCP classification policy contains entries used to map traffic received on access-uplink ports, or traffic received in the context of a network IP interface configured on a network port or hybrid port.

The **no** form of the policy disables the use of this policy.

Default

no dscp-classification-policy

Parameters

policy-id

Specifies the policy on the 7210 SAS.

Values 1 to 65535

mpls-lsp-exp-classification

Syntax

mpls-lsp-exp-classification *policy-id*
no mpls-lsp-exp-classification

Context

config>qos>network>ingress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates an mpls-lsp-exp classification policy to a forwarding class and profile state based on the MPLS EXP bits in the MPLS packet. The mpls-lsp-exp classification policy contains entries used to map traffic received in the context of a network IP interface configured on a network port or hybrid port.

The **no** form of this policy disables the use of this policy.

Default

no mpls-lsp-exp-classification

Parameters

policy-id

Specifies the policy on the 7210 SAS.

Values 1 to 65535

fc

Syntax

fc *fc-name* [create]
no fc *fc-name*

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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 queue mappings for *fc-name* forwarding types. The queue 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 fc-name

Specifies a case-sensitive, system-defined forwarding class name.

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.

use-dei

Syntax

[no] **use-dei**

Context

config>qos>network>ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables DEI based classification. When enabled for packets classified to this FC, the DEI bit is used to determine the ingress profile for the packet. Packets received with DEI bit set to zero are treated as in-profile and packets with DEI bit set to one are treated as out-of-profile packets.

When DEI based classification is enabled under the FC context, it overrides the profile values specified in the classification entry used to assign the FC.

The **no** form of this command disables use of DEI bit for classification of packets.

Default

no use-dei

7.7.2.1.6 Network ingress queue QoS policy commands

queue

Syntax

queue *queue-id* [**create**]

no queue *queue-id*

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies a queue for which to modify queue parameters.

The queue ID to FC map is user defined. In other words, the user can map FC to queues identified by queue IDs as per their needs.

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

Parameters

queue-id

Specifies the ID of the queue.

Values 1 to 8

create

Keyword to create a network queue policy.

adaptation-rule

Syntax

adaptation-rule [**cir** *adaptation-rule*] [**pir** *adaptation-rule*]

no adaptation-rule

Context

config>qos>network>ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

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 **rate** and **cir** 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.

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.

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.

cbs

Syntax

cbs *size-in-kbytes*

no cbs

Context

config>qos>network>ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the CBS value (minimum depth of the queue).

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

Default

32

Parameters

size-in-kbytes

Specifies the CBS value, in kilobytes.

Values	0 to 10240 (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T)
	0 to 102400 (7210 SAS-K 3SFP+ 8C)

mbs

Syntax

mbs *size-in-kbytes*

no mbs

Context

config>qos>network>ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the MBS value (maximum depth of the queue).
The **no** form of this command reverts to the default value.

Default

512

Parameters

size-in-kbytes

Specifies the MBS value, in kilobytes.

Values 0 to 12800 (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T)
 0 to 63488 (7210 SAS-K 3SFP+ 8C)

priority

Syntax

priority *level*
no priority

Context

config>qos>network>ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the queue priority. The queue priority is used by the scheduler to determine the order of service in both the within-cir loop and within-pir loop. Higher priority queues are serviced before lower priority queues.
The **no** form of this command reverts to the default value.

Default

1

Parameters

level

Specifies the priority of the queue.

Values 1 to 4

rate

Syntax

rate [*cir cir-percent*] [*pir pir-percent*]

no rate

Context

config>qos>network>ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

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 **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(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 is assumed. The parameter must be given as 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 is assumed. The parameter must be given as positive integer.

Values 1 to 100

Default 100

slope-policy

Syntax

slope-policy *name*

no slope-policy

Context

config>qos>network>ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default slope-policy configuration for the queue. The specified slope policy name must exist as a current slope policy name. If the slope policy does not exist, the **slope-policy** command fails. If a slope policy is currently associated with a queue, the slope policy cannot be removed from the system.

The slope policy contains the ring and non-ring high and low WRED slope definitions that are used by the queue. See [Buffer pools](#) for more information about ring and non-ring buffer pools and slope usage.

If the **slope-policy** command is not executed or the **no slope-policy** command is executed, the default slope policy is associated with the queue.

The **no** form of this command reverts the queue to the default slope policy.

Parameters

name

Specifies an existing slope policy name, up to 32 characters.

weight

Syntax

weight *weight*

no weight

Context

config>qos>network>ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the weight of the queue.

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.

Default

1

Parameters

weight

Specifies the weight of the queue.

Values 1 to 100

7.7.2.1.7 Self-generated traffic commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

sgt-qos

Syntax

sgt-qos

Context

config>router

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure DSCP and dot1p marking values 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
```

```
config>service>vprn>sgt-qos
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures DSCP and dot1p marking values for self-generated application traffic. When an application is configured using this command, the specified DSCP name or value is used for all packets generated by this application within the router instance in which it is configured. The instances can be base router or VPRN service.

The values configured in this command are used to perform the following actions:

- The values are used to set the DSCP bits in the IP packet.
- The values are mapped to the FC.
- The values are used to set the Ethernet 802.1p and MPLS EXP bits in the egress QoS policy, as based on the FC. This includes ARP and IS-IS packets that, because of their nature, do not carry DSCP bits.
- The values are used to configure the IP DSCP value in the packet IP header when the remarking command is disabled in the egress QoS policy associated with the service object from which the packet is sent out of. If an egress QoS policy is configured and IP DSCP or dot1p remarking is enabled, the DSCP or dot1p bits in the packet header are remarked based on the FC assigned to the packet.

Only one DSCP name or value can be configured per application. If multiple entries are configured, 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, dhcp, dns, ftp, icmp, igmp, ldp, ndis, ntp, ospf, pim, ptp, radius,
rip, rsvp, snmp, snmp-notification, ssh, syslog, tacplus, telnet, tftp,
traceroute, vrrp

The following values apply to the VPRN instance:

bgp, dhcp, icmp, igmp, ndis, ospf, pim, ssh, telnet, traceroute, vrrp

dscp-value

Specifies a value when this packet egresses. The respective egress policy should provide the mapping for the DSCP value to either LSP-EXP bits or dot1p bits as appropriate; otherwise, 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, 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

config>service>vprn>sgt-qos

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates a mapping between the DSCP of the self-generated traffic and the forwarding class.

Self-generated traffic for configured applications that matches the specified DSCP is 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-FC association.

Parameters

dscp-name

Specifies the name of the DSCP to be associated with the forwarding class. Only an existing DSCP can be specified, and it can only be specified by its name. 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 use the configured IP DSCP value.

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

7.7.2.2 Show commands

network

Syntax

network [*network-policy-id*] **association**

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

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays network policy information.

Parameters

network-policy-id

Displays network information for the specific policy ID.

Default all network policies

Values 1 to 65535

detail

Displays information about ingress and egress dot1p EXP bit mappings and network policy interface associations.

association

Displays policy associations.

Output

The following output is an example of QoS network policy information, and [Table 37: Output fields: QoS network](#) describes the output fields.

Sample output

```
*A:Dut-A>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       : None
Scope          : Template
DOT1P Class Poli*: 1                      DSCP Class Polic*: 1
MPLS Lsp Exp Cla*: 1
Description    : Default network QoS policy.
-----
FC      Queue      MCast Queue Use Dei
-----
be      1          1          false
l2      2          2          false
af      3          3          false
l1      4          4          false
h2      5          5          false
ef      6          6          false
h1      7          7          false
nc      8          8          false
-----
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           100      closest
-----
Queue Priority and Weight Details
-----
-----
QueueId      Priority      Weight
-----
Queue1       1            1
Queue2       1            1
Queue3       1            1
Queue4       1            1
Queue5       1            1
Queue6       1            1
Queue7       1            1
Queue8       1            1
-----
High Slope Non Ring
-----
-----
```


QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	80
Queue2	Down	70	90	80
Queue3	Down	70	90	80
Queue4	Down	70	90	80
Queue5	Down	70	90	80
Queue6	Down	70	90	80
Queue7	Down	70	90	80
Queue8	Down	70	90	80
Low Slope Non Ring				
QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	80
Queue2	Down	50	75	80
Queue3	Down	50	75	80
Queue4	Down	50	75	80
Queue5	Down	50	75	80
Queue6	Down	50	75	80
Queue7	Down	50	75	80
Queue8	Down	50	75	80
High Slope Ring				
QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	80
Queue2	Down	70	90	80
Queue3	Down	70	90	80
Queue4	Down	70	90	80
Queue5	Down	70	90	80
Queue6	Down	70	90	80
Queue7	Down	70	90	80
Queue8	Down	70	90	80
Low Slope Ring				
QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	80
Queue2	Down	50	75	80
Queue3	Down	50	75	80
Queue4	Down	50	75	80
Queue5	Down	50	75	80
Queue6	Down	50	75	80
Queue7	Down	50	75	80
Queue8	Down	50	75	80
Slope Policies				
QueueId	CBS(KBytes)	MBS(KBytes)	Slope-Policy	
Queue1	24	200	default	
Queue2	24	200	default	
Queue3	24	200	default	
Queue4	24	200	default	
Queue5	24	200	default	
Queue6	24	200	default	

```

Queue7      24      200      default
Queue8      24      200      default
-----
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
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A>show>qos#

*A:SAH01-051>show>qos# network 1 association
=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id      : 1
Egr Remark     : False
Forward Class  : be                      Profile      : None
Scope          : Template
DOT1P Class Po*: 1                      DSCP Class P*: 0
Description    : Default network-port QoS policy.

-----
Port Attachments
-----
No Matching Entries

=====
* indicates that the corresponding row element may have been truncated.
*A:SAH01-051>show>qos#

```

Table 37: Output fields: QoS network

Label	Description
Policy-Id	The ID that uniquely identifies the 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 — Indicates that packets are classified as out-profile In — Indicates packets are classified as in-profile None — Indicates packets profile is undefined
DOT1P Class	Specifies the dot1p-classification policy ID that is being used for mapping the packets to different FC under the FCs based on the dot1p bits

Label	Description
DSCP classification	Specifies the dscp-classification policy ID that is being used for mapping the packets to different FC under the FCs based on the dscp bits
High Slope Non Ring	Specifies the non-ring high-slope policy values
Low Slope Non Ring	Specifies the non-ring low-slope values
High Slope Ring	Specifies the ring high-slope values
Slope Policies	Displays the slope policies applied to the queues

router

Syntax

router [*router-instance*]

router service-name *service-name*

Context

show

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command shows router information.

Parameters

router-instance

Specifies the router name or service ID.

Values *router-name* — Base
 service-id — 1 to 2147483647

Default Base

service-name

Specifies the service name, up to 64 characters.

sgt-qos

Syntax

sgt-qos

Context

show>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays QoS information about self-generated traffic. The value "none" in the output indicates that the default value is used; it does not mean that there is no value set. For a list of application defaults, see [Table 32: Network policy ID #1 defaults for access-uplink ports](#).

application

Syntax

application [app-name] [dscp | dot1p]

Context

show>router>sgt-qos

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays application QoS settings.

Parameters

app-name

Specifies the application name.

- Values**
- The following values apply to the base router instance:

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

The following values apply to a VPRN service instance:

bgp, icmp, igmp, ndis, ospf, pim, ssh, telnet, traceroute, vrrp

**Note:**

- The **ptp** value in the context of SGT QoS is defined as Precision Timing Protocol (PTP) and is an application. The PTP application name is also used in areas like event-control and logging. PTP is defined in IEEE 1588-2008.
- The **ptp** value in the context of IP filters is defined as Performance Transparency Protocol (PTP). 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 are examples of application QoS information, and [Table 38: Output fields: SGT-QoS application](#) describes the output fields.

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
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 38: 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

Supported on all 7210 SAS platforms as described in this document

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 output is an example of DSCP-to-FC mapping information, and [Table 39: Output fields: SGT-QoS DSCP map](#) describes the output fields.

Output Sample for 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#
```

Output Sample (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 39: Output fields: SGT-QoS DSCP map

Label	Description
DSCP Value	The DSCP values (displayed as names) of the self-generated traffic
FC Value	The FC value mapped to each DSCP value
Default FC Value	The default FC value

8 Network queue QoS policies

This chapter provides information to configure network queue QoS policies using the CLI.

8.1 Overview

On the 7210 SAS-K 2F1C2T, network queue policies define the egress network queuing for the traffic egressing on the access-uplink ports. Network queue policies are used at the Ethernet port and define the bandwidth distribution for the FC traffic egressing on the Ethernet port. The user can define the number of queues and the mapping of FC-to-queue per network-queue policy. Each of these queues are shared by unicast and multicast traffic.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, network queue policies define the egress network queuing for the traffic egressing on the access-uplink port, and for the traffic egressing from the network IP interface configured on the network port or hybrid port. Network queue policies define the bandwidth distribution for the FC traffic egressing on the Ethernet port. The user can define the number of queues and the mapping of FC-to-queue per network-queue policy. Each of these queues are shared by unicast and multicast traffic.

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 Creating a network queue QoS policy

Configuring and applying QoS policies other than the default policy is optional. A default network queue policy is applied to all access uplink ports, network ports, and hybrid ports.

To create a network queue policy, define the following:

- **network queue policy name**

The system does not dynamically assign a name.

- **description**

The description provides a brief overview of policy features.

The FC-to-queue ID mapping can be defined by the user in the policy. The user has an option to use fewer queues.

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]
```

Example

```
*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#
```

8.2.2 Applying network queue policies

This section provides information about applying network queue policies.

8.2.2.1 Applying network queue configuration to access-uplink ports

Use the following syntax to apply a network queue policy to an Ethernet port in access-uplink port mode.

```
config>port#
    ethernet
        network
            queue-policy policy-name
```

Example

```
#-----
echo "Port Configuration"
#-----
    port 1/1/1
        ethernet
            access
                uplink
                    queue-policy "nql-cbs"
            exit
    exit
no shutdown
    exit
#-----
```

8.2.2.2 Applying network queue configuration to network ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Use the following syntax to apply a network queue policy to an Ethernet network port in network port mode or hybrid port mode.

```
config>port#
    ethernet
        network
            queue-policy policy-name
```

Example

```
#-----
echo "Port Configuration"
#-----
    port 1/1/1
        ethernet
            network
                queue-policy "nql-cbs"
            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.

8.3.1 Default network queue policy for 7210 SAS-K 2F1C2T

Example

The following is a sample of default policy parameters for the 7210 SAS-K 2F1C2T.

```
*A:dut-i>config>qos>network-queue# info detail
-----
description "Default network queue QoS policy."
queue 1 create
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 2 create
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 3 create
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 4 create
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 5 create
  rate cir 100 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 6 create
  rate cir 100 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
```

```
exit
queue 7 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 50
    slope-policy "default"
    priority 1
    weight 1
exit
queue 8 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 50
    slope-policy "default"
    priority 1
    weight 1
exit
fc af create
    queue 3
exit
fc be create
    queue 1
exit
fc ef create
    queue 6
exit
fc h1 create
    queue 7
exit
fc h2 create
    queue 5
exit
fc l1 create
    queue 4
exit
fc l2 create
    queue 2
exit
fc nc create
    queue 8
exit
-----
*A:dut-i>config>qos>network-queue#
```

8.3.2 Default network queue policy for 7210 SAS-K 2F6C4T

Example

The following is a sample of default policy parameters for the 7210 SAS-K 2F6C4T.

```
*A:K-SASK12>config>qos>network-queue# info detail
-----
description "Default network queue QoS policy."
queue 1 create
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
```

```
        priority 1
        weight 1
    exit
    queue 2 create
        rate cir 25 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 3 create
        rate cir 25 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 4 create
        rate cir 25 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 5 create
        rate cir 100 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 6 create
        rate cir 100 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 7 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 8 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
```



```
        weight 1
    exit
    fc af create
        queue 3
    exit
    fc be create
        queue 1
    exit
    fc ef create
        queue 6
    exit
    fc h1 create
        queue 7
    exit
    fc h2 create
        queue 5
    exit
    fc l1 create
        queue 4
    exit
    fc l2 create
        queue 2
    exit
    fc nc create
        queue 8
    exit
-----
*A:K-SASK12>config>qos>network-queue#
```

8.3.3 Default network queue policy for 7210 SAS-K 3SFP+ 8C

Example

The following are sample default policy parameters for the 7210 SAS-K 3SFP+ 8C.

```
*A:Dut-A>config>qos>network-queue# info detail
-----
description "Default network queue QoS policy."
queue 1 create
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
    priority 1
    weight 1
exit
queue 2 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
    priority 1
    weight 1
exit
queue 3 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
```

```
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 4 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 5 create
        rate cir 15 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 6 create
        rate cir 15 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 7 create
        rate cir 5 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    queue 8 create
        rate cir 10 pir 100
        adaptation-rule cir closest pir closest
        mbs 200
        cbs 24
        slope-policy "default"
        priority 1
        weight 1
    exit
    fc af create
        queue 3
    exit
    fc be create
        queue 1
    exit
    fc ef create
        queue 6
    exit
    fc h1 create
        queue 7
    exit
    fc h2 create
        queue 5
    exit
```

```
fc l1 create
    queue 4
exit
fc l2 create
    queue 2
exit
fc nc create
    queue 8
exit
-----
*A:Dut-A>config>qos>network-queue#
```

8.4 Service management tasks

This section describes the service management tasks.

8.4.1 Deleting network queue QoS policies

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

Example

The following shows the command usage to delete a user-created network queue policy.

```
config>qos# no network-queue policy-name
```

Example

The following example shows the command usage to delete a user-created network queue policy.

```
config>qos# no network-queue nq1
```

8.4.2 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: Command usage to copy and overwrite a QoS policy

```
config>qos# copy network-queue nq1-cbs nq2-cbs
```

Example: Configuration output showing copied policies

```
*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
-----
*A:card-1>config>qos# info
```

8.4.3 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.

8.5 Network queue QoS policy command reference

- [Command hierarchies](#)
- [Command descriptions](#)

8.5.1 Command hierarchies

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

8.5.1.1 Configuration commands

```
- config
  - qos
    - network-queue policy-name [create]
      - description description-string
      - no description
      - [no] queue queue-id
        - adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
        - no adaptation-rule
        - cbs size-in-kbyte
        - no cbs
        - mbs size in kbytes
        - no mbs
        - priority level
        - no priority
        - slope-policy name
        - no slope-policy
        - rate cir cir-percent [pir pir-percent]
        - no rate
        - weight weight
        - no weight
```

8.5.1.2 Operational commands

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

8.5.1.3 Show commands

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

8.5.2 Command descriptions

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

8.5.2.1 Configuration commands

- [Generic commands](#)
- [Operational commands](#)
- [Network queue QoS policy commands](#)
- [Network queue QoS policy queue commands](#)

8.5.2.1.1 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

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.

8.5.2.1.2 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

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 attempts to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy. This parameter indicated that the source policy ID and the destination policy ID are network-queue policy IDs.

overwrite

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

8.5.2.1.3 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

Description

This command configures a network queue policy. Network queue policies on the Ethernet port define network egress queuing.

On the 7210 SAS-K 2F1C2T, the network queue policy can be associated with access-uplink ports to define queues to be used on egress.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the network queue policy can be associated with access-uplink ports and network ports to define queues to be used on egress.

Default

default

Parameters

policy-name

Specifies 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.

8.5.2.1.4 Network queue QoS policy queue commands

queue

Syntax

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

Context

config>qos>network-queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a QoS network-queue policy queue.

The user has an option to define the FC to queue map. Either one or multiple FCs can be mapped to the same queue.

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

Parameters

queue-id

Specifies the *queue-id* for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy.

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

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 **rate** and **cir** 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.

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.

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.

cbs

Syntax

[no] **cbs** *size-in-kbytes*

Context

config>qos>network-queue>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the CBS value (minimum depth of the queue).

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

Default

32

Parameters

size-in-kbytes

Specifies the CBS value in kilobytes.

Values 0 to 10240 (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
 0 to 102400 (7210 SAS-K 3SFP+ 8C)

mbs

Syntax

mbs *size-in-kbytes*
no mbs

Context

config>qos>network-queue>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the MBS value (maximum depth of the queue).
The **no** form of this command reverts to the default value.

Default

512

Parameters

size-in-kbytes

Specifies the MBS value in kilobytes.

Values 0 to 12800 (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
 0 to 63488 (7210 SAS-K 3SFP+ 8C)

priority

Syntax

priority *level*
no priority

Context

config>qos>network-queue>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the queue priority. The queue priority is used by the scheduler to determine the order of service in both the within-cir loop and within-pir loop. Higher priority queues are serviced before lower priority queues.

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

Default

1

Parameters

level

Specifies the priority of the queue.

Values 1 to 4

slope-policy

Syntax

slope-policy *name*

no slope-policy

Context

config>qos>network-queue>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default slope-policy configuration for the queue. The specified slope-policy name must exist as a current slope policy name. If the slope policy does not exist, the **slope-policy** command fails. If a slope policy is currently associated with a queue, the slope policy cannot be removed from the system.

The slope policy contains the ring and non-ring high and low WRED slope definitions that are used by the queue. See [Buffer pools](#) for more information about ring and non-ring buffer pools and slope usage.

If the **slope-policy** command is not executed or the **no slope-policy** command is executed, the default slope policy will be associated with the queue.

The **no** form of this command reverts the queue to the default slope policy.

Parameters

name

Specifies an existing slope policy name, up to 32 characters.

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

Description

This command defines the administrative PIR and the administrative CIR parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets 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 **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(0) parameters.

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 is assumed. The parameter must be given as positive integer.

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

Default 0

Values 0 to 100

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 is assumed. The parameter must be given as positive integer.

Default	100
Values	1 to 100

weight

Syntax
`weight` *weight*
`no weight`

Context
config>qos>network-queue>queue

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command specifies the weight of the queue.
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 default.

Parameters
weight
Specifies the weight of the queue.
Values 1 to 100

8.5.2.2 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

Description

This command displays network queue policy information.

Parameters

- network-queue-policy-name**
Displays network queue information for the specified 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 each queue's rates and adaptation-rule and cbs details. It also shows FC to queue mapping details.

Output

The following output is an example of network queue information, and [Table 40: Output fields: network queue](#) describes the output fields.

Sample output

```
*A:SAH01-051>show>qos# network-queue "default" detail

=====
QoS Network Queue Policy
=====
-----
Network Queue Policy (default)
-----
Policy          : default
Description     : Default network queue QoS policy.
Wrr Policy      :
Pkt.Byte Offset: 0
-----
-----
FC To Queue Mappings
-----
FC      Queue
-----
be      1
l2      2
af      3
l1      4
h2      5
ef      6
h1      7
nc      8
-----
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	100	closest

Queue Priority and Weight Details				

QueueId	Priority	Weight		

Queue1	1	1		
Queue2	1	1		
Queue3	1	1		
Queue4	1	1		
Queue5	1	1		
Queue6	1	1		
Queue7	1	1		
Queue8	1	1		

High Slope Non Ring				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	70	90	80
Queue2	Down	70	90	80
Queue3	Down	70	90	80
Queue4	Down	70	90	80
Queue5	Down	70	90	80
Queue6	Down	70	90	80
Queue7	Down	70	90	80
Queue8	Down	70	90	80

Low Slope Non Ring				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	50	75	80
Queue2	Down	50	75	80
Queue3	Down	50	75	80
Queue4	Down	50	75	80
Queue5	Down	50	75	80
Queue6	Down	50	75	80
Queue7	Down	50	75	80
Queue8	Down	50	75	80

High Slope Ring				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	70	90	80
Queue2	Down	70	90	80
Queue3	Down	70	90	80
Queue4	Down	70	90	80
Queue5	Down	70	90	80
Queue6	Down	70	90	80
Queue7	Down	70	90	80
Queue8	Down	70	90	80

Low Slope Ring				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)


```

-----
Queue1      Down      50      75      80
Queue2      Down      50      75      80
Queue3      Down      50      75      80
Queue4      Down      50      75      80
Queue5      Down      50      75      80
Queue6      Down      50      75      80
Queue7      Down      50      75      80
Queue8      Down      50      75      80
-----

```

Slope Policies

```

-----
QueueId      CBS(KBytes)  MBS(KBytes)  Slope-Policy
-----
Queue1       50          200          default
Queue2       50          200          default
Queue3       50          200          default
Queue4       50          200          default
Queue5       50          200          default
Queue6       50          200          default
Queue7       50          200          default
Queue8       50          200          default
-----

```

Network-Port Associations

```

-----
No Matching Entries
-----

```

```

=====
*A:SAH01-051>show>qos#
=====

```

Table 40: Output fields: network queue

Label	Description
Policy	The policy name that uniquely identifies the policy
Accounting	Displays whether the accounting mode is packet-based or frame-based
Description	A text string that helps identify the policy's context in the configuration file
Port-Id	Displays the physical port identifier where the network queue policy is applied
Queue	Displays the queue ID
CIR	Displays the committed information rate
PIR	Displays the peak information rate
CBS	Displays the committed burst size
FC	Displays FC to queue mapping

9 Service ingress QoS policies

This chapter provides information to configure SAP ingress QoS policies using the CLI.

9.1 Overview of service ingress policy

There is one default service ingress policy. The default policy allocates a single queue and maps all traffic to the "be" (best-effort) FC. The default policies can be copied and modified but they cannot be deleted. The default policies are identified as policy ID 1. The default policies are applied to the appropriate interface, by default. For example, the default SAP-ingress policy is applied to access ingress SAPs. You must explicitly associate other QoS policies. For information about the tasks and commands necessary to access the command line interface and to configure and maintain your 7210 SAS devices, see the "CLI Usage" chapter in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide*.

In a service ingress QoS policy, a user can define up to eight queues or 16 policers/meters per policy, with up to two queues or two meters per FC.

For VPLS, the following types of forwarding are supported:

- unicast
- multicast
- broadcast
- 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. All these traffic types use the same queue or meter (in other words, a separate queue or meter for multicast, broadcast, and unknown unicast traffic types cannot be defined).

Unicast and multipoint traffic can be defined to use the same queue or meter or different queues or meter per FC. In other words, eight queues and 16 policers are shared by unicast and multicast traffic types and if a user allocates a dedicated multicast queue for BUM traffic, the number of queues available for unicast traffic reduces. Similarly, up to 16 meters can be shared by unicast and multicast traffic types.

9.1.1 Configuration guidelines for SAP-ingress policy

The configuration guidelines for SAP ingress policies are the following:

- The option is available to configure a smaller number of ingress queues or ingress policers per SAP.
- On the 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T, hardware queues are allocated in groups of two, and on the 7210 SAS-K 3SFP+ 8C, hardware queues are allocated in groups of four; these grouped queues are reserved for use by the SAP even if the user specifies an odd value. The 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T support two, four, six, or eight queues per SAP. The 7210 SAS-K 3SFP+ 8C supports four or eight queues per SAP.

- FC-to-queue map can be defined; this allows the user to assign the packets classified into a particular FC to any one of the queues configured for the SAP.
- The option is available to use up to two queues per FC, with one queue for unicast traffic and one queue for BUM traffic, with a maximum of eight queues per access SAP. This option can be used with multipoint services for example, VPLS service. BUM traffic shares a single queue per FC; therefore it is not possible to use individual queue for each of broadcast, unknown-unicast, and multicast traffic. It is possible to define the same queue for unicast and BUM traffic. For example, users can assign two queues per FC, such that unicast traffic uses one of the queues and the BUM traffic uses the other queue. This allows users to have four FCs per SAP with two queues per FC; or the user can have seven FCs per SAP with one queue per FC and the eighth queue being shared by BUM traffic of all the FCs; or a mix and match is allowed. If a multicast queue is not assigned to an FC explicitly, it uses queue 1 (the default queue of the policy).
- The queue parameters such as queue shaper rate (CIR/PIR), CBS and MBS, queue priority, and weight can be defined. The assigned priority and weight are used to determine the priority and weight of the queue in both the CIR and PIR scheduling loop.
- Allow configuration of WRED slopes (per queue) – high-slope and low-slope. Depending on the queue mode and the profile assigned to the packet on SAP ingress classification, one of the configured WRED slopes is used to evaluate if a buffer can be allocated to the packet. In-profile packets use the high-slope and out-of-profile packets use the low-slope.
- Supports Strict priority (SP) scheduling and Weighted-Fair Queuing (WFQ) scheduling for SAP ingress queues. For more information about service ingress scheduling, see the [Schedulers](#).
- The option is available to use up to two policers per FC, with one policer for unicast traffic and one for BUM traffic, with a maximum of up to 16 policers per access SAP. This option can be used with multipoint services for example, VPLS service. BUM traffic shares a single meter per FC; therefore it is not possible to use individual meter for each of broadcast, unknown-unicast, and multicast traffic. It is possible to define the same meter for unicast and BUM traffic. For example, users can assign two meters per FC, such that unicast traffic uses one meter and the BUM traffic uses the other meter. This allows users to have eight FCs per SAP with two meters per FC, or users can have eight FCs per SAP with one meter per FC and the eighth meter being shared by BUM traffic of all the FCs; or a mix and match is allowed. If a multicast meter is not assigned to an FC explicitly, it uses queue 1 (the default queue of the policy).
- The option is available to use a queue or a meter per FC, with up to two queues per FC or two meters per FC or a queue and a meter per FC for a maximum of up to eight queues per access SAP and 16 meters/policers per access SAP. It is allowed to configure queue for unicast traffic and meter for BUM traffic or the other way around. This option can be used with multipoint services for example, VPLS service. BUM traffic shares a single meter or queue per FC; therefore it is not possible to use individual meter or an individual queue for each of broadcast, unknown-unicast, and multicast traffic. If a multicast meter is not assigned to an FC explicitly, it uses queue 1 (the default queue of the policy).
- The meter parameters such as meter rate (CIR/PIR), CBS and MBS, and meter mode (srTCM, trTCM) can be defined.
- If both IP criteria and MAC criteria are configured in SAP ingress QoS classification then they need resources from two different slices. In other words, the MAC and IP criteria entries cannot be located in the same slice. If resources are not found in two different slices, then the association of the policy fails.

9.1.1.1 Resource allocation for service ingress QoS classification policy

The available global pool of ingress internal CAM hardware resources can be allocated as per user needs for use with different features such as SAP ingress QoS policy, ingress ACLs, and so on. SAP ingress QoS can be allocated classification for use from this pool. Resources can be allocated for SAP ingress QoS policy classification IPv4, IPv6, and MAC match criteria, based on the operator needs. Users can modify the resources allocated to scale the number of entries available per match criteria or scale the number of SAPs. The resources from the global ingress internal CAM pool are allocated in slices with a fixed number of entries.

The number of slices to be allotted for a SAP ingress QoS policy is specified using the **config system resource-profile ingress-internal-tcam qos-sap-ingress-resource** CLI command.

The user can specify a limit for the amount of resources required for SAP ingress QoS policies and also has the option to limit the amount of resources used per match criteria supported for SAP ingress QoS policies. A specific slice can be used for MAC criteria, IP criteria, and IPv6 criteria.

Before associating SAP-ingress policy match criteria with a SAP, resources must be allocated. Until resources are allocated for use, attempts to associate a policy with a SAP fail. When the user allocates resources for use by SAP ingress QoS policies using the **config system resource-profile ingress-internal-tcam qos-sap-ingress-resource** CLI command, the system allocates resources in slices of 510 entries (192 entries for the 7210 SAS-K 3SFP+ 8C).

The above resources set the maximum limit on the resources available for use by all SAP ingress policies in use simultaneously on the system. The software manages the resource slices allocated to the SAP ingress QoS policy pool and allocates the entries in the slices when a SAP ingress QoS policy is associated with a SAP. The software allocates the resources required by a SAP from the slices depending on whether the SAP-ingress policy uses IP criteria, IPv6 criteria, or MAC criteria, and the number of entries configured in the SAP-ingress policy.

When the user allocates slices of resources using the **config system resource-profile ingress-internal-tcam qos-sap-ingress-resource** CLI command, the resources are used only for classification entries configured under IPv4 criteria or MAC criteria.

To use IPv6 criteria, the user must allocate resources using the **config system resource-profile ingress-internal-tcam qos-sap-ingress-resource mac-ipv4-ipv6-128-match-enable** CLI command. These resources are used for classification entries configured under the IPv6 criteria, IPv4 criteria, or MAC criteria. Each IPv6, IPv4, or MAC classification entry consumes two resources from this pool, reducing the number of classification entries that can be accommodated in a single slice to 255 (96 entries for the 7210 SAS-K 3SFP+ 8C.) The user can choose to allocate all the slices or a portion from the allocation for SAP ingress QoS classification for IPv6 criteria.

The **tools dump system-resources** CLI command displays the current usage and availability of resources. One or more entries per slice are reserved for system use.

9.1.1.2 Resource allocation for SAP ingress meters

The FC ingress meter and SAP ingress aggregate meter allocate resources for the common meter resource pool. A single FC ingress meter requires one entry from the pool when the policy is associated with the SAP and a single SAP ingress aggregate meter uses a single entry from the pool when the command to enable SAP aggregate meter functionality is executed under the context of the SAP. An increase in FC ingress meter reduces the number of meters available for SAP ingress aggregate meter. The reverse is also true.

The **tools dump system-resources** CLI command can be used to display the current utilization of the ingress meter resource pool.

9.1.1.3 Default SAP-ingress policy

The default policy 1 maps all traffic to default FC "be" and maps FC "be" to queue 1. Queue 1 is configured with CIR 0 and PIR max.

Example: Configuration output of a default SAP-ingress policy

```
A:SAH01-071>config>qos>sap-ingress# info detail
-----
policy-name "default"
description "Default SAP ingress QoS policy."
no ip-mac-match
scope template
queue 1 create
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
  slope-policy "default"
  mbs 60
  cbs 10
  priority 1
  weight 1
exit
fc "af" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "be" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "ef" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "h1" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "h2" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "l1" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "l2" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "nc" create
```

```

queue 1
multicast-queue 1
no use-dei
exit
default-fc "be"
no dot1p-classification
no dscp-classification

```

The following table lists the SAP-ingress policy defaults.

Table 41: SAP-ingress policy defaults

Field	Default
description	"Default SAP-ingress QoS policy."
scope	template
queue	1
adaptation-rule	cir closest pir closest
rate	pir = max, cir= 0
cbs	10KBytes - default
mbs	60KBytes - default
priority	1
weight	1
default-fc	be

9.1.1.4 Use of index file by SAP QoS ingress policy

The 7210 SAS uses an index file to store the map that indicates the QoS resource allocation to the SAPs. This file is used to ensure that all the SAPs that were created successfully before a reboot can be recreated during a reboot. 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 index file is stored in the flash. During a reboot, if the file is found, the system allocates resources as per 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. When the index file is not present it is possible that the saved configuration did not execute successfully after the reboot.



Note:

The following restrictions apply:

- There is no guarantee that resources will be allocated to all SAPs.
- The index file used for QoS maps is different from the one used for storing interface indexes.

9.2 Basic configurations

A basic service ingress QoS policy must conform to the following:

- have a unique service ingress QoS policy ID
- have a QoS policy scope of **template** or **exclusive**
- have at least one default unicast FC queue
- (optionally) use multipoint FC queue
- (optionally) use unicast meter and BUM traffic

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP, a default QoS policy is applied.

9.2.1 Service ingress QoS policies

About this task

To create a service ingress policy, perform the following.

Procedure

- Step 1.** Define a policy ID value.
The system does not dynamically assign a value.
- Step 2.** Define a description that provides a brief overview of policy features.
- Step 3.** Specify a default FC for the policy.
All packets received on an ingress SAP using this ingress QoS policy are classified to the default FC.
- Step 4.** Define FC parameters by performing the following.
- Modify the unicast/queue default value to override the default unicast forwarding type queue mapping for **fc fc-name**.
 - Modify the multicast or queue default value to override the default multicast forwarding type queue mapping for **fc fc-name**.
 - Associate a meter for **fc fc-name**, if required, for both unicast and BUM traffic types.
- Step 5.** Specify the following 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 FC for matched traffic.
- Step 6.** Create a SAP-ingress policy using 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.

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"
...
-----
A:ALA-7>config>qos>sap-ingress#
```

9.2.2 Service ingress QoS queues

About this task

To create service ingress queue parameters, perform the following.

Procedure

- Step 1.** Define a new queue ID value.
The system does not dynamically assign a value.
- Step 2.** Configure queue parameters. Rate, slope-policy, CBS, MBS, priority, and weight.
The following is a sample ingress queue configuration output.

Example

```
A:ALA-7>config>qos# info
echo "QoS Policy Configuration"
#-----
    sap-ingress 1 create
    policy-name "default"
    description "Default SAP ingress QoS policy."
    no ip-mac-match
    scope template
    queue 1 create
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    slope-policy "default"
    mbs 60
    cbs 10
    priority 1
    weight 1
    exit
.....
#-----
A:ALA-7>config>qos#
```

9.2.3 Service ingress QoS meters

Example

The following are sample service ingress QoS meters.

```
A:ALA-7>config>qos#
echo "QoS Policy Configuration"
#-----
sap-ingress 1 create
policy-name "test"
description "SAP ingress QoS policy."
no ip-mac-match
scope template
```



```
meter 1 create
adaptation-rule cir closest pir closest
rate cir 0 pir max
mbs 60
cbs 10
exit
.....
#-----
A:ALA-7>config>qos#
```

9.2.4 SAP ingress FC configuration

Example

The following is a sample SAP ingress FC configuration output with unicast queues and multicast queues.

```
*A:dut-i>config>qos>sap-ingress$ info
-----
....
queue 1 create
exit
queue 2 create
rate cir 1000 pir 2000
exit
fc "af" create
queue 1
multicast-queue 2
exit
fc "nc" create
queue 2
multicast-queue 1
exit
.....
-----
config>qos>sap-ingress$ info
```

Example

The following is a sample SAP ingress FC configuration output with a mix of unicast queues and meters, and a mix of multicast queues and meters.

```
*A:dut-i>config>qos>sap-ingress$ info
-----
....
queue 1 create
exit
queue 2 create
rate cir 1000 pir 2000
exit
meter 1 create
rate cir 10000 pir 10000
exit
meter 2 create
rate cir 1000 pir 2000
exit
fc "h2" create
meter 1
multicast-meter 2
exit
```

```
fc "nc" create
meter1
multicast-meter 2
exit
fc "af" create
queue 1
multicast-queue 2
exit
fc "be" create
queue 2
multicast-queue 1
exit
.....
-----
config>qos>sap-ingress$ info
```

9.2.5 Service ingress dot1p and IP DSCP criteria

Example

The following is a sample configuration output of dot1p classification policy and IP DSCP classification policy used for ingress classification and its association with SAP-ingress policy.

```
A:ALA-7>config>qos>dot1p-classification# info
#-----
.....
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
.....

#-----

A:ALA-7>config>qos>sap-ingress# info
#-----
.....

dot1p-classification 1

.....
#-----
A:ALA-7>
```

9.2.6 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.

Example: Ingress IP criteria configuration output

```
A:ALA-7>config>qos# info
...
```

```
#-----
echo "QoS Policy Configuration"
#-----
...
        sap-ingress 100 create
...
        ip-criteria
        entry 10 create
        description "Entry 10-FC-AF"
        match dscp af12
        exit
        action fc af
        exit
        entry 20 create
        description "Entry 20-FC-BE"
        match dscp be
        exit
        no action
        exit
        exit
exit
..
#-----
A:ALA-7>config>qos#
```

9.2.7 Service ingress MAC match criteria

About this task

To configure service ingress policy MAC criteria, perform the following:

Procedure

- Step 1.** Define a new entry ID value.
Entries must be explicitly created; the system does not dynamically assign entries or a value.
- Step 2.** Associate the FC with a specific MAC criteria entry ID.
- Step 3.** Define a description.
The description provides a brief overview of policy features.

Example

The following is a sample ingress MAC criteria configuration output.

```
A:ALA-7>config>qos# info
...
#-----
echo "QoS Policy Configuration"
#-----
...
        sap-ingress 101 create
...
        mac-criteria
        entry 10 create
        description "Entry10"
        match
            dst-mac 04-67-ff-00-00-01 ff-ff-ff-ff-ff-ff
            dot1p 7 7
        exit
        action fc be
```

```
        exit
      exit
    exit
#-----
A:ALA-7>config>qos#
```

9.2.8 Applying service ingress policies

SAP ingress QoS policies are supported only on access SAPs.

9.2.8.1 Epipe service

Example

The following sample configuration output shows an Epipe service configuration with SAP-ingress policy 100 applied to the SAP.

```
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
      exit
    exit
  exit
exit
-----
A:ALA-7>config>service#
```

9.2.8.2 VPLS

Example

The following sample configuration output shows a VPLS service configuration with SAP-ingress policy 100.

```
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
      exit
    exit
  exit
exit
-----
A:ALA-7>config>service#
```

9.2.8.3 IES

Example

The following sample configuration output shows an IES service configuration.

```
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
      exit
    exit
  exit
  no shutdown
exit
...
-----
```

9.3 Service management tasks

This section describes service management tasks.

9.3.1 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 it is applied.

Example

```
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#
```

9.3.1.1 Removing a QoS policy from service SAPs

Example

The following Epipe service output sample shows that the SAP service ingress reverts to *policy-id* 1 when the non-default policies are removed from the configuration.

```
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#
```

9.3.2 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.

```
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
```

9.3.3 Removing 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
```

9.3.4 Editing QoS policies

You can change existing QoS 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 overwrite the original policy.

9.4 Service SAP QoS policy command reference

- [Command hierarchies](#)
- [Command descriptions](#)

9.4.1 Command hierarchies

- [Service ingress QoS policy commands](#)
- [Operational commands](#)
- [Show commands](#)

9.4.1.1 Service ingress QoS policy commands

```
- config
- qos
- [no] sap-ingress policy-id [create]
  - default-fc fc [profile {in |out | use-dei}]
  - 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]
    - multicast-meter meter-id
    - no multicast-meter
    - no meter
    - meter meter-id
    - multicast-queue queue-id
    - no multicast-queue
```

```
- no queue
- queue queue-id
- no use-dei
- use-dei
- [no] ip-criteria
  - [no] entry entry-id [create]
    - action [fc fc-name] [profile {in | out | use-dei}]
    - no action
    - description description-string
    - no description
    - match [protocol protocol-id]
    - no match
      - dscp dscp-name
      - no dscp
      - dst-ip {ip-address/mask | ip-address netmask}
      - no dst-ip
      - dst-port fc {eq} dst-port-number
      - no dst-port
      - fragment {true | false}
      - no fragment
      - src-ip {ip-address/mask | ip-address netmask}
      - no src-ip
      - src-port {eq} src-port-number
      - no src-port
    - dscp [old-entry-id new-entry-id]
  - [no] ipv6-criteria
    - [no] entry entry-id [create]
      - action [fc fc-name] [profile {in | out | use-dei}]
      - no action
      - description description-string
      - no description
      - match [next-header next-header]
      - no match
        - dscp dscp-name
        - no dscp
        - dst-ip {ip-address/mask | ip-address netmask}
        - no dst-ip
        - dst-port fc {eq} dst-port-number
        - no dst-port
        - fragment {true | false}
        - no fragment
        - src-ip {ip-address/mask | ip-address netmask}
        - no src-ip
        - src-port {eq} src-port-number
        - no src-port
      - dscp [old-entry-id new-entry-id]
  - [no] mac-criteria
    - [no] entry entry-id [create]
      - action [fc fc-name] [profile {in | out | use-dei}]
      - no action
      - description description-string
      - no description
      - [no] match
        - dst-mac ieee-address [ieee-address-mask]
        - no dst-mac
        - etype 0x0600..0xffff
        - no etype
        - 0inner-dot1p dot1p-value [dot1p-mask]
        - no 0inner-dot1p
        - inner-tag value [vid-mask]
        - no inner-tag
        - no outer-dot1p
        - outer-dot1p dot1p-value [dot1p-mask]
        - no outer-tag
```



```

- outer-tag value [vid-mask]
- src-mac ieee-address [ieee-address-mask]
- no src-mac
- dscp old-entry-id new-entry-id
- queue queue-id
- no queue
- [no] adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
- cbs size-in-kbyte
- no cbs
- mbs size in kbytes
- no mbs
- no priority
- priority level
- no rate
- rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
- no slope-policy
- slope-policy name
- no weight
- weight weight
- no meter meter-id
- [no] adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
- cbs size-in-kbyte
- no cbs
- mbs size in kbytes
- no mbs
- mode mode
- no mode
- rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
- no rate
- scope {exclusive | template}
- no scope

```

9.4.1.2 Operational commands

```

- config
- qos
- copy sap-ingress src-pol dst-pol [overwrite]

```

9.4.1.3 Show commands

```

- show
- qos
- sap-ingress policy-id [detail | association | match-criteria]
- dot1p-classification [policy-id] [detail]
- dscp-classification [policy-id] [detail]

```

9.4.2 Command descriptions

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

9.4.2.1 Configuration commands



Note:

The 7210 SAS QoS capability varies across platforms. In the command descriptions, the term queue/meter is used, and depending on the platform capabilities, either one or both terms can be used.

9.4.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>sap-ingress

config>qos>sap-ingress>ip-criteria>entry

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

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.

9.4.2.2 Operational commands

copy

Syntax

```
copy sap-ingress src-pol dst-pol [overwrite]  
copy dot1p-classification src-pol dst-pol [overwrite]  
copy dscp-classification src-pol dst-pol [overwrite]
```

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

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 that the source policy ID and the destination policy ID are SAP ingress policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

Values 1 to 65535

overwrite

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

dot1p-classification *src-pol-dst-pol*

Specifies that the source policy ID and the destination policy ID are dot1p classification policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

Values 1 to 65535

dscp-classification *src-pol-dst-pol*

Specifies that the source policy ID and the destination policy ID are IP DSCP classification policy IDs. Specify the source policy ID that the copy command will attempt to copy from

and specify the destination policy ID to which the command will copy a duplicate of the policy.

Values 1 to 65535

dscp

Syntax

dscp *dscp-name*

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

Description

This command configures a DiffServ Code Point (DSCP) code point to be used as a network ingress QoS policy match criterion.

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

Parameters

dscp-name

Specifies a dscp name that has been previously mapped to a value using the dscp-name command. The DiffServ code point can only be specified by its name.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

dst-ip

Syntax

dst-ip {*ip-address/mask* | *ip-address netmask*}

dst-ip {*ipv6-address/prefix-length*}

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

Description

This command configures a destination address range to be used as a network 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 be used only for IPv4.

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-address	a.b.c.d
ipv6-address	x::x::x::x::x::x (eight 16-bit pieces)
	x::x::x::x::d.d.d.d
	x — 0 to FFFF (hexadecimal)
	d — 0 to 255 (decimal)

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

Description

This command configures a destination TCP or UDP port number for a network ingress 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)

fragment

Syntax

fragment {true | false}

no fragment

Context

config>qos>ingress>ip-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures fragmented or non-fragmented IP packets as a network ingress QoS policy match criterion.

The **no** form of this command removes the match criterion and matches all packets regardless of whether they are fragmented or not.

Parameters

true

Keyword to configure a match on all fragmented IP packets. A match occurs 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.

src-ip

Syntax

src-ip {*ip-address/mask* | *ip-address ipv4-address-mask*}

src-ip {*ipv6-address/prefix-length*}

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

Description

This command configures a source IPv4 or IPv6 address range to be used as a network ingress QoS policy match criterion.

To match on the source IPv4 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 for IPv4.

The **no** form of this command removes the source IPv4 or IPv6 address match criterion.

Default

no source IP match criterion

Parameters

ip-address

Specifies the source IPv4 address specified in dotted decimal notation.

Values *ipv4-address* — a.b.c.d

mask

Specifies the length in bits of the subnet mask.

Values 1 to 32

ipv4-address-mask

Specifies the subnet mask, in dotted decimal notation.

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

ipv6-address

Specifies the IPv6 prefix for the IP match criterion, in hexadecimal digits.

Values	ipv6-address	x::x::x::x::x::x (eight 16-bit pieces)
	—	
		x::x::x::x::d.d.d.d
		x — 0 to FFFF (hexadecimal)
		d — 0 to 255 (decimal)

prefix

Specifies the IPv6 prefix length for the ipv6-address, expressed as a decimal integer.

Values 1 to 128

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

Description

This command configures a source TCP or UDP port number for a network ingress 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

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

Description

This command rennumbers existing QoS policy criteria entries to properly sequence policy entries.

This is 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 number of an existing entry.

Values 1 to 64

new-entry-id

Specifies the new entry number to be assigned to the old entry.

Values 1 to 64

9.4.2.3 Service ingress QoS policy commands

sap-ingress

Syntax

```
[no] sap-ingress policy-id [create]
```

Context

```
config>qos
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates or edits the ingress policy. The ingress policy defines the Service Level Agreement (SLA) enforcement that service packets receive as they ingress a SAP; the SAP can be configured on an access port or hybrid port. SLA enforcement is accomplished through the definition of meters or 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/meter that all ingress traffic flows through. Complex policies have multiple meters/queues combined with classification entries that indicate which meter/queue a packet flows 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.

Depending on the support available on different 7210 SAS platforms, SAP ingress policies can be defined with either dot1p as the match criteria or IP DSCP as the match criteria or IP headers as the match criteria or MAC headers or both as the match criteria. On the 7210 SAS, the user has an option to use dot1p, IP DSCP, IPv4/IPv6 criteria and MAC criteria.

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 about the default SAP Ingress policy for different platforms.

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:

- A SAP ingress policy with multiple criteria can be associated with the SAP. The system defines a match-order to match against the classification rules specified in the policy. Please see the QoS overview section to know more about the match order.
- Before using IPv4, IPv6 and MAC match criteria, resources must be allocated using the CLI command **config>system>resource-profile>ingress-internal-tcam>qos-sap-ingress-resource**. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for information about resource allocation. See [Service ingress QoS policies](#) for more information about this CLI command.
- Service ingress only queues with shapers are 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 restores the factory defaults to policy ID 1.

Parameters

policy-id

Specifies the policy.

Values 1 to 65535

create

Keyword to create a sap ingress policy.

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 associates a dot1p classification policy, which contains entries used to map traffic received on a hybrid port or SAP to a forwarding class and profile state, based on the dot1p bits in the packet.

When it is defined in a service ingress QoS policy and associated with an access SAP ingress, a packet received on the SAP is used to match with the dot1p values defined in this policy. If a match is found, the corresponding forwarding class and profile are assigned to the packet.

The **no** form of this command disables use of dot1p classification policy.

Parameters

policy-id

Specifies the policy.

Values 1 to 65535

dscp-classification

Syntax

dscp-classification

no dscp-classification

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command associates an IP DSCP classification policy, which contains entries used to map traffic received on a SAP to a forwarding class and profile state, based on the IP DSCP bits in the packet.

When it is defined in a service ingress QoS policy and associated with an access SAP ingress, a packet received on the SAP is used to match with the IP DSCP values defined in this policy. If a match is found, the corresponding forwarding class and profile are assigned to the packet.

The **no** form of this command disables use of DSCP classification policy.

Parameters

policy-id

Specifies the policy ID that uniquely identifies the policy.

Values 1 to 65535

scope

Syntax

scope {exclusive | template}

no scope

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

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.

default-fc

Syntax

default-fc *fc* [**profile** {**in** | **out** | **use-dei**}]

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

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 is 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 are 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/meter. The value for *fc* must be one of the predefined forwarding classes in the system.

profile {**in** | **out**}

Specifies that packets matching the classification entry are explicitly classified to either in-profile or out-of-profile. To remove the profile action for a classification entry, the **action** command must be re-executed without the profile action defined. The profile assigned by the user is used subsequently to determine the slope to use at the ingress and egress queuing points and is used for egress marking (if enabled).

in

Specifies that any packets matching the classification rule are treated as in-profile. Mutually exclusive to the **out** parameter following the profile classification action keyword.

out

Specifies that any packets matching the classification rule are treated as out-of-profile. Mutually exclusive to the **in** parameter following the profile classification action keyword.

use-dei

Syntax

[no] use-dei

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables DEI based classification. When enabled, the packet classified to this FC the DEI bit is used to determine the ingress profile for the packet. Packets received with DEI bit set to zero are treated as in-profile and packets with DEI bit set to one are treated as out-of-profile packets.

When DEI based classification is enabled, it overrides the profile values specified in the classification entry used to assign the FC.

The **no** form of this command disables use of DEI bit for classification of packets.

Default

no use-dei

fc

Syntax

[no] fc *fc-name* [create]

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a class instance of the forwarding class FC name. After 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 of 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.

multicast-queue

Syntax

multicast-queue *queue-id*

no multicast-queue

Context

config>qos>sap-ingress>fc

config>qos>network>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default multicast forwarding type queue mapping for **fc** *fc-name*. The specified queue ID must exist within the policy as a multipoint queue before the mapping can be made. After the forwarding class mapping is executed, all broadcast, unknown unicast, and multicast (BUM) traffic on a SAP or a access-uplink port using this policy is forwarded using the queue ID.

The **no** form of this command removes the association of the FC and the queue. When the no form is executed, the BUM traffic uses the default multicast forwarding type queue.

Default

1

Parameters

queue-id

Specifies the queue ID.

Values 1 to 8

ip-criteria

Syntax

[no] ip-criteria

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context create or edit policy entries that specify IP criteria used to match supported fields from the IP packet header, including IP DSCP.

IP criteria-based SAP ingress policies are used to select the 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.



Note:

Before associating a SAP ingress policy configured to use IPv4 criteria with a SAP, resources must be allocated to it. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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. When IP criteria entries are removed from a SAP ingress policy, the IP criteria is removed from all services where that policy is applied.

ipv6-criteria

Syntax

[no] ipv6-criteria

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context create or edit policy entries that specify IPv6 criteria such as IP quintuple lookup or DiffServ code point.

IPv6 criteria-based SAP ingress policies are used to select the forwarding class for matched traffic.

The 7210 SAS implementation exits on the first match found and execute the actions in accordance with the accompanying **action** command. For this reason entries must be sequenced correctly from most to least explicit.



Note:

Before associating a SAP ingress policy configured to use IPv6 criteria with a SAP, resources must be allocated to it. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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. When ipv6-criteria entries are removed from a SAP ingress policy, the ipv6-criteria is removed from all services where that policy is applied.

mac-criteria

Syntax

[no] mac-criteria

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context create or edit policy entries that specify MAC criteria.

The MAC criteria-based SAP ingress policies are used to select the forwarding class for matched traffic.

The 7210 SAS implementation exits on the first match found and execute the actions in accordance with the accompanying **action** command. For this reason entries must be sequenced correctly from most to least explicit.



Note:

Before associating a SAP ingress policy configured to use MAC criteria with a SAP, resources must be allocated to it. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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. When mac-criteria entries are removed from a SAP ingress policy, the mac-criteria is removed from all services where that policy is applied.

queue

Syntax

[no] queue queue-id

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default unicast forwarding type queue mapping for **fc** *fc-name*. The specified queue ID must exist within the policy as a non-multipoint queue before the mapping can be made. When the forwarding class mapping is executed, all unicast traffic (this includes all traffic, even broadcast and multicast for services) on a SAP or an access-uplink port using this policy is forwarded using the queue ID.

The **no** form of this command sets the unicast (point-to-point) queue ID back to the default queue for the forwarding class (queue 1).

Parameters

queue-id

Specifies the queue ID for the queue.

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

Description

This command overrides the default multicast forwarding type queue mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy before the mapping can be made. When the forwarding class mapping is executed, all broadcast, unknown unicast, and multicast (BUM) traffic on an access SAP using this policy is forwarded using the meter ID.

The default multicast forwarding type queue is the default unicast queue (queue 1). In other words, if no multicast meter is assigned to a FC, it uses default queue 1.

The **no** form of this command removes the association of the FC and the meter. When the no form is executed, the BUM traffic uses the default multicast forwarding type queue (queue 1).

Default

1

Parameters

meter-id

Specifies the meter ID for the meter.

Values 1 to 16

meter

Syntax

[no] **meter** *meter-id*

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default unicast forwarding type queue mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy before the mapping can be made. When the forwarding class mapping is executed, all unicast traffic (this includes all traffic, even broadcast and multicast for services) on an access SAP using this policy is forwarded using the meter ID.

The **no** form of this command sets the unicast (point-to-point) traffic type back to the default queue for the forwarding class (queue 1).

Parameters

meter-id

Specifies the meter ID for the meter.

Values 1 to 16

9.4.2.4 Service ingress QoS policy entry commands

action

Syntax

action [**fc** *fc-name*] **profile** {in|out}

no action

Context

config>qos>sap-ingress>ip-criteria>entry

```
config>qos>sap-ingress>mac-criteria>entry
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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 fc-name** 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](#) command

Parameters

fc fc-name

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

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

profile {in | out}

Specifies the profile assignment. When specified, packets matching the classification entry are explicitly classified to either in-profile or out-of-profile. The remove the profile action for an classification entry, the **action** command must be re-executed without the profile action defined. The profile assigned by the user is used subsequently to determine the slope to use at the ingress and egress queuing points and is used for egress marking (if enabled).

in

Specifies that any packets matching the classification rule are treated as in-profile. The **in** parameter is mutually exclusive to the **out** parameter following the profile classification **action** keyword. Either **in** or **out** must be specified when the profile keyword is present.

out

Specifies that any packets matching the classification rule are treated as out-of-profile. The **out** parameter is mutually exclusive to the **in** parameter following the profile classification action keyword. Either **in** or **out** must be specified when the profile keyword is present.

use-dei

Specifies whether DEI must be used to determine the initial profile of the packet.

entry

Syntax

[no] **entry** *entry-id* [**create**]

Context

config>qos>sap-ingress>ip-criteria

config>qos>sap-ingress>mac-criteria

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context create or edit 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 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 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 are considered incomplete and therefore are rendered inactive.

Values 1 to 64

create

Keyword to create a flow entry when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit

the object. This keyword is not required when the protection is disabled. The keyword is ignored when the flow entry already exists.

match

Syntax

[no] match [**protocol** *protocol-id*]

Context

config>qos>sap-ingress>ip-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a context to configure match criteria for SAP ingress 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 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

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; check the node for values.

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

Description

Commands in this context create and edit match MAC criteria for ingress SAP QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured then all criteria must be satisfied (and function) before the action associated with the match 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*.

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

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 meader to match.

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

Values protocol numbers accepted in DHB: 0 to 42, 45 to 49, 52 to 59, 61 to 255 keywords: none, crtp, crudp, egg, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp * — udp/tcp wildcard

9.4.2.5 Service ingress MAC QoS policy match commands

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

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	0xFFFFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

All packets with a source MAC OUI value of 00-03-FA subject to a match condition should be specified as: 0003FA000000 0x0FFFFFF000000

Default 0xFFFFFFFFFFFFFF (hex) (exact match)

Values 0x0000000000000000 to 0xFFFFFFFFFFFFFF (hex)

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

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 type II frame Ether type value to be used as a match criterion, expressed in hexadecimal or decimal notation.

Values 0x0600 to 0xFFFF (hexadecimal)
1536 to 65535 (decimal)

0inner-dot1p

Syntax

inner-dot1p *dot1p-value* [*dot1p-mask*]

no inner-dot1p

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the dot1p value to be used as a service ingress QoS policy match criterion to match against the dot1p value in the inner tag (the one that follows the outermost tag in the packet) of the received packet.

The **no** form of this command removes the previously entered dot1p value as the match criteria.

Default

no inner-dot1p

Parameters

dot1p-value

Specifies the dot1p value to match.

Values 0 to 7

dot1p-mask

Specifies the mask value to match a range of dot1p values, expressed in hexadecimal or decimal notation.

Values 0 to 7

inner-tag

Syntax

inner-tag *value* [*vid-mask*]

no inner-tag

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the VLAN value to be used as a service ingress QoS policy match criterion to match against the VLAN value in the inner tag (the one that follows the outermost tag in the packet) of the received packet.

The **no** form of this command removes the previously entered VLAN tag value as the match criteria.

Default

no inner-tag

Parameters

value

Specifies the VLAN value to use for the match.

Values 0 to 4095 (decimal)
0x0 to 0xFFFF (hexadecimal)

vid-mask

Specifies the mask value to match a range of VLAN values.

Values 1 to 4095 (decimal)
1x0 to 0xFFFF (hexadecimal)

outer-dot1p

Syntax

outer-dot1p *dot1p-value* [*dot1p-mask*]

no outer-dot1p

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the dot1p value to be used as a service ingress QoS policy match criterion to match against the dot1p value in the outermost tag of the received packet.

The **no** form of this command removes the previously entered dot1p value as the match criteria.

Default

no outer-dot1p

Parameters

dot1p-value

Specifies the dot1p value to match.

Values 0 to 7

dot1p-mask

Specifies the mask value to match a range of dot1p values, expressed in hexadecimal or decimal notation.

Values 0 to 7

outer-tag

Syntax

outer-tag *value* [*vid-mask*]

no outer-tag

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the VLAN value to be used as a service ingress QoS policy match criterion to match against the VLAN value in the outermost tag in the packet of the received packet.

The **no** form of this command removes the previously entered VLAN tag value as the match criteria.

Default

no outer-tag

Parameters

value

Specifies the VLAN value to use for the match.

Values 0 to 4095 (decimal)
 0x0 to 0xFF (hexadecimal)

vid-mask

Specifies the mask value to match a range of VLAN values.

Values 0 to 4095 (decimal)
 0x0 to 0xFF (hexadecimal)

src-mac

Syntax

src-mac *ieee-address* [*ieee-address-mask*]
no src-mac

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

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 address 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	DDDDDDDDDDDDDDDD	281474959933440

Format style	Format syntax	Example
Hexadecimal	0xHHHHHHHHHHHHH	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

Default0xFFFFFFFFFFFFFF (hexadecimal) (exact match)

Values0x0000000000000000 to 0xFFFFFFFFFFFFFF (hexadecimal)

9.4.2.6 SAP ingress queue and meter QoS policy commands

queue

Syntax

queue queue-id [create]

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context create a queue and modify queue parameters associated with a particular queue. The queue-ID to FC map is user-defined.

The 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T support two, four, six, or eight queues per SAP. The 7210 SAS-K 3SFP+ 8C supports four or eight queues per SAP

The **no** form of this command removes the queue.

Parameters

queue-id

Specifies the ID of the queue.

Values1 to 8

create

Keyword to create a network queue policy.

meter

Syntax

meter *meter-id*

no meter

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command modifies meter parameters associated with a particular meter. The queue-ID to FC map is user-defined.

The **no** form of this command remove the meter definition.

Default

no meter

Parameters

meter-id

Specifies the ID of the meter.

Values 1 to 16

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]

no adaptation-rule

Context

config>qos>sap-ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

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 cir and pir apply.

Default

adaptation-rule pir closest cir 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.

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.

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.

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]

no adaptation-rule

Context

config>qos>sap-ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

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.

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.

cbs

Syntax

[no] **cbs** *size-in-kbytes*

Context

config>qos>sap-ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the CBS value (minimum depth of the queue).

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

Default

10

Parameters

size-in-kbytes

Specifies the CBS in kilobytes.

Values 0 to 10240 (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
 0 to 102400 (7210 SAS-K 3SFP+ 8C)

cbs

Syntax

[no] **cbs** *size* [*kbits* | *bytes* | *kbytes*]

Context

config>qos>sap-ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the CBS value of 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 out-profile by the meter to indicate that the traffic is complying with the meter-configured CIR rate.

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

Default

kbits

Parameters

[kbits | bytes | kbytes]

Specifies the maximum burst size.

Values size-in-kbits — 1 to 16384, default
 size-in-bytes — 64 to 2097152, default
 size-in-kbytes — 1 to 2048, default

mbs

Syntax

[no] **mbs** *size-in-kbytes*

Context

config>qos>sap-ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the MBS value (maximum depth of the queue).

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

Default

60

Parameters

size-in-kbytes

Specifies the MBS value in kilobytes.

Values	0 to 12800 (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
	0 to 63488 (7210 SAS-K 3SFP+ 8C)

mbs

Syntax

[no] mbs size [*kbits* | *bytes* | *kbytes*]

Context

config>qos>sap-ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the MBS value of the meter. The maximum burst size parameter specifies the maximum burst size that can be transmitted by the source while still 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 complying with the meter-configured PIR rate.

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

Default

kbits

Parameters

[kbits | bytes | kbytes]

Specifies the maximum burst size.

Values	size-in-kbits — 1 to 16384, default
---------------	-------------------------------------

size-in-bytes — 64 to 2097152, default
size-in-kbytes — 1 to 2048, default

priority

Syntax
[no] priority *level*

Context
config>qos>sap-ingress>queue

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command defines the queue priority. The queue priority is used by the scheduler to determine the order of service in both the within-cir loop and within-pir loop. Higher priority queues are serviced before lower priority queues.
The **no** form of this command reverts to the default value.

Default
1

Parameters
level
Specifies the priority of the queue.
Values 1 to 4

mode

Syntax
mode *mode*
no mode

Context
config>qos>sap-ingress>meter

Platforms
Supported on all 7210 SAS platforms as described in this document

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 sets it to default mode.

Default

trtcm2

Parameters

mode

Specifies the mode of the meter.

Values srtcm — 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 Maximum Burst Size (MBS). A packet is marked green if it does not exceed the CBS and yellow if it does exceed the CBS, but not the MBS; otherwise, it is marked red and is dropped. The srTCM mode is useful, for example, for ingress policing of a service, where only the length, not the peak rate, of the burst determines service eligibility.

trtcm2 — implements the policing algorithm defined in RFC 4115. The trTCM2 mode meters the packet stream and marks the packets either green, yellow, or red. A packet is marked red and is dropped 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 mode is useful, for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate.



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.

rate

Syntax

rate **cir** *cir-rate-in-kbps* [**pir** *pir-rate-in-kbps*]

no rate

Context

config>qos>sap-ingress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

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 `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.

The `no` form of this command returns all queues created with the queue ID by association with the QoS policy to the default PIR and CIR parameters (`max`, 0).

Default

`rate 0 pir max`

Parameters

cir-rate-in-kbps

Specifies the administrative CIR rate, in kilobits, for the queue. If the `rate` command is not executed or the `cir` parameter is not explicitly specified, the default CIR value is used.

Default	0
Values	0 to 3000000, max (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T) 0 to 20000000, max (7210 SAS-K 3SFP+ 8C)

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.

Default	max
Values	1 to 3000000, max (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T) 1 to 20000000, max (7210 SAS-K 3SFP+ 8C)

rate

Syntax

`rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]`
`no rate`

Context

`config>qos>sap-ingress>meter`

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the meter. The PIR defines the maximum rate at which the meter can admit the packet into the system for forwarding. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the flow can be limited by oversubscription factors or available egress bandwidth.

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.

The **no** form of this command returns all meters created with the *meter-id* by association with the QoS policy to the default PIR and CIR parameters (max, 0).

Default

rate 0 pir max

Parameters

cir-rate-in-kbps

Specifies the administrative CIR rate, in kilobits, for the queue. If the rate command is not executed or the cir parameter is not explicitly specified, the default CIR value is used.

Default	0
Values	0 to 3000000, max (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T) 0 to 20000000, max (7210 SAS-K 3SFP+ 8C)

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.



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.

Default	max
Values	1 to 3000000, max (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T) 1 to 20000000, max (7210 SAS-K 3SFP+ 8C)

9.4.2.7 Show commands

```
sap-ingress
```

Syntax

```
sap-ingress [policy-id] [detail | association | match-criteria]
```

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.

Default all SAP ingress policies

Values 1 to 65535

detail

Displays detailed policy information including policy associations.

associations

Displays the policy associations of the sap-ingress policy.

match-criterion

Displays the match-criterion of the sap-ingress policy.

Output

The following output is an example of SAP ingress QoS policy information, and [Table 42: Output fields: SAP ingress](#) describes the output fields.

Sample output

```
*A:Dut-B>show>qos# sap-ingress 1 detail
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (1)
-----
Policy-id           : 1           Scope           : Template
Default FC         : be
```

```

Criteria-type           : None
Mac Sub-Criteria       : None
IPv6 Enabled           : False
DOT1P Class Policy Id  : 0
MPLS Lsp Exp Class Policy*: 0
Name                   : default
Description             : Default SAP ingress QoS policy.

```

FC	Queue	MCast Queue Use Dei
be	1	1 false
l2	1	1 false
af	1	1 false
l1	1	1 false
h2	1	1 false
ef	1	1 false
h1	1	1 false
nc	1	1 false

FC	Meter	MCast Meter Use Dei
be	-	- false
l2	-	- false
af	-	- false
l1	-	- false
h2	-	- false
ef	-	- false
h1	-	- false
nc	-	- false

Queue Rates and Rules

QueueId	CIR	CIR Adpt Rule	PIR	PIR Adpt Rule
Queue1	0	closest	max	closest

Queue Priority and Weight Details

QueueId	Priority	Weight
Queue1	1	1

High Slope Non Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	80

Low Slope Non Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	80

High Slope Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	80

```

-----
Low Slope Ring
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down        50            75          80
-----
Slope Policies
-----
QueueId      CBS(KBytes)  MBS(KBytes)   Slope-Policy
-----
Queue1       10          60            default
-----
Match Criteria
-----
No Matching Criteria.
SAP Associations
-----
No Associations Found.
-----
Meter Mode      CIR Admin CIR Rule PIR Admin PIR Rule CBS Admin MBS Admin
                CIR Oper          PIR Oper          CBS Oper  MBS Oper
-----
No Matching Entries
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-B>show>qos#

```

Table 42: Output fields: SAP ingress

Label	Description
Policy-Id	The ID that uniquely identifies the policy
Scope	Exclusive — Implies that this policy can only be applied to a single SAP Template — Implies that this policy can be applied to multiple SAPs on the router
Description	A text string that helps identify the policy's context in the configuration file
Default FC	Specifies the default forwarding class for the policy
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
Meter	Displays the meter ID
Mode	Specifies the configured mode of the meter (trTCM2 or srTCM)

Label	Description
CIR Admin	Specifies the administrative Committed Information Rate (CIR) parameters for the meters
CIR Rule	min — The operational CIR for the meters will be equal to or greater than the administrative rate specified using the rate command max — The operational CIR for the meter will be equal to or less than the administrative rate specified using the rate command 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
PIR Admin	Specifies the administrative Peak Information Rate (PIR) parameters for the meters
PIR Rule	min — The operational PIR for the meter will be equal to or greater than the administrative rate specified using the rate command max — The operational PIR for the meters will be equal to or less than the administrative rate specified using the rate command closest — The operational PIR for the meters will be the rate closest to the rate specified using the rate command
CBS	def — Specifies the default CBS value for the meters value — Specifies the value to override the default reserved buffers for the meters
MBS	def — Specifies the default MBS value value — Specifies the value to override the default MBS for the meter
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
Match Criteria	Specifies an IP or MAC criteria entry for the policy
DSCP	Specifies a DiffServ Code Point (DSCP) name used for an ingress SAP QoS policy match
FC	Specifies the entry's forwarding class

Label	Description
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 Ether type value to be used as a Service Ingress QoS policy match
FC	Specifies the entry's forwarding class
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
Classifiers required	Indicates the number of classifiers for a VPLS or Epipe service
Meters required	Indicates the number of meters for a VPLS or Epipe service
Meters mode	Displays the configured meter mode

dot1p-classification

Syntax

dot1p-classification [*policy-id*] [**detail**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays dot1p-classification QoS policy information.

Parameters

policy-id

Displays information about the specific policy ID.

detail

Displays detailed policy information.

Output

The following output is an example of dot1p classification policy information, and [Table 43: Output fields: dot1p classification](#) describes the output fields.

Sample output

```
A:SAH01-050>config>qos>sap-ingress# show qos dot1p-classification 10
=====
DOT1P Classification Maps
=====
-----
Dot1P Class Id      : 10
Description         : (Not Specified)
-----
A:SAH01-050>config>qos>sap-ingress#

A:SAH01-050>config>qos>sap-ingress# show qos dot1p-classification 10 detail

=====
DOT1P Classification Maps
=====
-----
Dot1P Class Id      : 10
Description         : (Not Specified)
-----
-----
Dot1P Bit Map      Forwarding Class      Profile
-----
6                   h1                   None
7                   nc                   None
-----
-----
Network Policy Associations
-----
No Network Policy Associations found.
-----
-----
SAP Ingress Associations
-----
SAP Ingress Id      : 10
-----
=====
```

Table 43: Output fields: dot1p classification

Label	Description
Dot1P Class Id	The ID that uniquely identifies the policy
Dot1P Bit Map	The dot1p value specified in the policy
Forwarding Class	The forwarding class to assigned to the packet if the received packet's dot1p values match the dot1p value configured

Label	Description
Description	A text string that helps identify the policy's context in the configuration file
Profile	Specifies the profile to be assigned to the packet

dscp-classification

Syntax

dscp-classification [*policy-id*] [**detail**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays DSCP classification QoS policy information.

Parameters

policy-id

Displays information about the specific policy ID.

detail

Displays detailed policy information.

Output

The following is an example of DSCP classification information, and [Table 44: Output fields: DSCP classification](#) describes the output fields.

Sample output

```
*A:SAH01-071>config>qos# show qos dscp-classification 6335 detail

=====
DSCP Classification Maps
=====
-----
Dscp Class Id      : 6335
Description        : (Not Specified)
-----

Dscp Bit Map      Forwarding Class      Profile
-----
cp7                nc                    In
cp15               h1                    In
cp23               ef                    In
cp31               h2                    In
```

cp39	l1	In
cp47	af	In
cp55	l2	In

Network Policy Associations		

No Network Policy Associations found.		

SAP Ingress Associations		

SAP Ingress Id	: 17	

=====		

Table 44: Output fields: DSCP classification

Label	Description
Dot1P Class Id	The ID that uniquely identifies the policy
Dot1P Bit Map	The dot1p value specified in the policy
Forwarding Class	The forwarding class to assigned to the packet if the received packet's dot1p values match the dot1p value configured
Description	A text string that helps identify the policy's context in the configuration file
Profile	Specifies the profile to be assigned to the packet

10 Service egress QoS policies

This chapter provides information to configure service egress QoS policies using the CLI.

10.1 Overview

The service-egress policy defines the Service Level Agreement (SLA) for service packets as they egress on a SAP configured on either an access port or a hybrid port. Service-egress QoS policies allow the definition of queue parameters along with a remark policy.

With the default service egress policy, the system allocates one queue. The eight FCs are mapped to use the same queue. The user has the option to define up to eight queues per policy and define the FC-to-queue mapping. In addition, the policy allows the user to define the queue parameters. The hardware does not support a linear range of values for the rate parameters (CIR and PIR). The user can specify the computation method of rates to match the rates supported by the hardware, through the configuration of adaptation rules.

The SAP egress policy for access SAPs supports the following:

- per-SAP egress queuing and shaping, hierarchical shaping on SAP egress (with three levels of shaping) with a per-FC/queue shaper, per-SAP aggregate shaper and per-port egress rate shaper
- SAP egress queues, shaping, and scheduling
 - On the 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T, hardware queues are allocated in groups of two, and on the 7210 SAS-K 3SFP+ 8C, hardware queues are allocated in groups of four; these grouped queues are reserved for use by the SAP even if the user specifies an odd value. The 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T support two, four, six, or eight queues per SAP. The 7210 SAS-K 3SFP+ 8C supports four or eight queues per SAP.
 - Provides an option to configure the FC-to-queue map, allowing the user to assign the packets classified into a particular FC to any of the queues configured for the SAP.
 - On SAP egress, only a single queue can be configured per FC and all traffic (unicast and BUM) share the single queue.
 - Allows configuration of queue shaper rate (CIR/PIR), CBS and MBS, queue priority and weight. The assigned priority and weight is used to determine the priority and weight of the queue in both the CIR and PIR scheduling loop.
 - Allows configuration of WRED slopes (per queue) – high-slope and low-slope. One of the configured WRED slopes is used to allocate buffer to the packet. In-profile packets use the high-slope and out-of-profile packets use the low-slope. The profile of the packet is determined at the ingress (access uplink port ingress or Access SAP ingress) and carried through to be used at SAP egress to determine the WRED slope to apply and also to determine the egress marking value to use (if remarking/marketing is enabled).
 - Supports Strict Priority (SP) scheduling and Weighted-Fair Queuing (WFQ) scheduling for SAP egress queues.
- SAP egress remarking/marketing

Dot1p and/or IP DSCP marking must be supported on access SAP egress. Configuration of per FC dot1p and/or IP DSCP marking is supported, with the capability to assign different dot1p and/or IP DSCP values for in-profile and out-of-profile packets. In addition, support for marking DEI value is available.

Example

The following is a sample configuration output showing the default access SAP egress QoS.

```
*A:SAH01-071>config>qos>sap-egress# info detail
-----
description "Default SAP egress QoS policy."
scope template
no remarking
remark 1
queue 1 create
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs 60
    cbs 10
    slope-policy "default"
    priority 1
    weight 1
exit
fc af create
    queue 1
exit
fc be create
    queue 1
exit
fc ef create
    queue 1
exit
fc h1 create
    queue 1
exit
fc h2 create
    queue 1
exit
fc l1 create
    queue 1
exit
fc l2 create
    queue 1
exit
fc nc create
    queue 1
exit
```

10.1.1 Egress SAP FC and profile overrides

The FC of an access egress packet can be changed to redirect the packet to an alternate queue than the ingress FC determination would have used. The profile of an access egress packet can also be changed to modify the congestion behavior within the egress queue. In both cases, egress marking decisions are based on the new FC and profile, instead of the FC or profile determined at ingress.

The SAP egress QoS policy supports the use of reclassification rules to override the ingress FC and profile of packets that egress a SAP where the QoS policy is applied.

dot1p, IP precedence, and DSCP entries can be defined, each with an explicit FC or profile override parameters. The reclassification logic for each entry follows the same basic hierarchical behavior as the classification rules within the SAP ingress QoS policy. IP DSCP entries have the highest match priority, followed by IP precedence and dot1p.

If the IP precedence values overlap with DSCP values by matching the same IP header TOC field, the DSCP entry parameters override or remove the IP precedence parameters. When none of the matched entries override a parameter, the ingress classification is preserved; that is, the FC and profile assigned by SAP/network ingress classification rules is carried over and used at egress.



Note:

The egress queue shaper does not reassign the profile for the packet. By default, the profile is assigned a value of undefined at access SAP egress when the profile is not defined explicitly by the user in the service egress policy reclassification entry. A value of undefined specifies that the profile assigned at ingress is carried over. Therefore, the user can use egress reclassification to assign only an FC without specifying a profile. The profile assigned at ingress by either the ingress classification entry or the ingress shaper is carried over. The egress classification entry modifies the profile only if the user explicitly configures the profile as in or out.

This capability is supported on access SAP egress for all services.

10.1.2 Configuration guidelines for access SAP egress policies

This section provides a list of configuration guidelines for access SAP egress policies.



Note:

This section applies to SAPs configured on access ports and hybrid ports.

The configuration guidelines for access SAP egress policies are the following:

- On the 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T, hardware queues are allocated in groups of two, and on the 7210 SAS-K 3SFP+ 8C, hardware queues are allocated in groups of four; these grouped queues are reserved for use by the SAP even if the user specifies an odd value. The 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T support two, four, six, or eight queues per SAP. The 7210 SAS-K 3SFP+ 8C supports four or eight queues per SAP.
- FC-to-queue map can be defined; this allows the user to assign the packets classified into a particular FC to any one of the queues configured for the SAP.
- Both unicast traffic and BUM traffic share a single queue per FC. In other words, unlike service ingress policy, it is not possible to assign different queues for BUM traffic and unicast traffic.
- The queue parameters such as queue shaper rate (CIR/PIR), CBS and MBS, queue priority and weight can be defined. The assigned priority and weight is used to determine the priority and weight of the queue in both the CIR and PIR scheduling loop.
- WRED slopes (per queue) can be configured: high-slope and low-slope. Depending on the queue mode and the profile assigned to the packet on SAP ingress classification, one of the configured WRED slopes is used to evaluate if a buffer can be allocated to the packet. In-profile packets use the high-slope and out-of-profile packets use the low-slope.
- SP scheduling and WFQ scheduling for SAP ingress queues are supported.

10.1.2.1 Basic configurations

A basic service egress QoS policy must conform to the following:

- have a unique service egress QoS policy ID
- have a QoS policy scope of **template** or **exclusive**
- have at least one FC queue

10.1.2.2 Creating an access SAP egress policy

To create an access 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:SAH01-051>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] fc           + Configure forwarding-class mappings
[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
```

Example

```
*A:SAH01-051>config>qos# info detail
sap-egress 1 create
  description "Default SAP egress QoS policy."
  scope template
  no remarking
  remark 1
  queue 1 create
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs 60
    cbs 10
    slope-policy "default"
    priority 1
    weight 1
  exit
  fc af create
    queue 1
  exit
  fc be create
    queue 1
  exit
  fc ef create
```

```
        queue 1
    exit
    fc h1 create
        queue 1
    exit
    fc h2 create
        queue 1
    exit
    fc l1 create
        queue 1
    exit
    fc l2 create
        queue 1
    exit
    fc nc create
        queue 1
    exit
exit
```

10.1.2.3 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.

10.2 Service egress policy command reference

- [Command hierarchies](#)
- [Command description](#)

10.2.1 Command hierarchies

- [Configuration commands](#)
- [Copy commands](#)
- [Show commands](#)

10.2.1.1 Configuration commands

```
- config
  - qos
    - sap-egress policy-id [create]
```

```
- [no] description description-string
- dot1p-classification policy-id
- no dot1p-classification
- dscp-classification policy-id
- no dscp-classification
- fc fc-name [create]
- no fc fc-name
- ip-prec-classification policy-id
- no ip-prec-classification
- queue queue-id [create]
- no queue queue-id
  - adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
  - no adaptation-rule
  - cbs size-in-kbyte
  - no cbs
  - mbs size in kbytes
  - no mbs
  - priority level
  - no priority
  - slope-policy name
  - no slope-policy
  - rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
  - no rate
  - weight weight
  - no weight
- scope {exclusive | template}
- remark policy-id
- no remark
- [no] remarking
```

10.2.1.2 Copy commands

```
- config
  - qos
    - copy sap-egress src-pol dst-pol [overwrite]
```

10.2.1.3 Show commands

```
- show
  - qos
    - sap-egress [policy-id] [detail | association]
```

10.2.2 Command description

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

10.2.2.1 Configuration commands

- [Generic commands](#)

- [SAP egress queue QoS policy commands](#)

10.2.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>sap-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

10.2.2.1.2 SAP egress queue QoS policy commands

sap-egress

Syntax

sap-egress *policy-id* [**create**]

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a SAP egress policy. The SAP egress policy determines the QoS treatment to packets at service egress.

When the policy is created, by default only one queue is created. The user can create up to 8 queues and associate it to different FCs on the SAPs to which this SAP egress policy is attached. The SAP egress policy allows the user to define the queue parameters for the eight queues.

Default

sap-egress 1

Parameters

policy-id

Specifies the SAP egress policy.

Values 1 to 65535

dot1p-classification

Syntax

dot1p-classification *policy-id*

no dot1p-classification

Context

config>qos>sap-egress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates a dot1p classification policy, which contains entries used to map traffic sent on an access SAP to an FC and profile state, based on the dot1p bits in the packet.

When the dot1p classification policy is defined in a service egress QoS policy and associated with an access SAP egress, a packet sent out of the SAP is used to match with the dot1p values defined in this policy. If a match is found, the corresponding FC and profile are assigned to the packet.



Note:

The FC and profile assigned by the dot1p entry match can be overridden by an IP precedence match or IP DSCP match. See [Egress SAP FC and profile overrides](#) for more information about the match hierarchy.

The **no** form of this command disables the use of the dot1p classification policy.

Default

no dot1p-classification

Parameters

policy-id

Specifies the policy ID.

Values 1 to 65535

dscp-classification

Syntax

dscp-classification

no dscp-classification

Context

config>qos>sap-egress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates an IP DSCP classification policy, which contains entries used to map traffic sent on a SAP to an FC and profile state, based on the IP DSCP bits in the packet.

When the DSCP classification policy is defined in a service egress QoS policy and associated with an access SAP egress, a packet sent out of the SAP is used to match with the IP DSCP values defined in this policy. If a match is found, the corresponding FC and profile are assigned to the packet.



Note:

The FC and profile assigned by the dot1p entry match or IP precedence entry match can be overridden by an IP DSCP match. See [Egress SAP FC and profile overrides](#) for more information about the match hierarchy.

The **no** form of this command disables the use of the DSCP classification policy.

Default

no dscp-classification

Parameters

policy-id

Specifies the policy ID.

Values 1 to 65535

fc

Syntax

fc *fc-name* [create]

no fc *fc-name*

Context

config>qos>sap-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a class instance of the forwarding class name. When 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.

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

create

Creates a forwarding class.

ip-prec-classification

Syntax

ip-prec-classification *policy-id*

no ip-prec-classification

Context

config>qos>sap-egress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates an IP DSCP classification policy with IP precedence entries used to map traffic sent out of a SAP to an FC and profile state based on the IP precedence bits in the packet.

When it is defined in a service egress QoS policy and associated with an access SAP egress, a packet sent out of the SAP is used to match with the IP precedence values defined in this policy. If a match is found, the corresponding FC and profile are assigned to the packet.



Note:

The FC and profile assigned by the dot1p entry match can be overridden by an IP precedence match, and the FC and profile assigned by the IP precedence match can be overridden by an IP DSCP match. See [Egress SAP FC and profile overrides](#) for more information about the match hierarchy.

The **no** form of this command disables the use of the DSCP classification policy and IP precedence values.

Default

no ip-prec-classification

Parameters

policy-id

Specifies the policy ID.

Values 1 to 65535

queue

Syntax

queue *queue-id* [create]

no queue *queue-id*

Context

config>qos>sap-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a queue and modifies queue parameters associated with a particular queue. The queue ID to FC map is user-defined.

The 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T support two, four, six, or eight queues per SAP.

The 7210 SAS-K 3SFP+ 8C supports four or eight queues per SAP.

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

Parameters

queue-id

Specifies the ID of the queue.

Values 1 to 8

create
Creates a network queue policy.

adaptation-rule

Syntax
adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]
no adaptation-rule

Context
config>qos>sap-egress>queue

Platforms
Supported on all 7210 SAS platforms as described in this document

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.

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.

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.

cbs

Syntax

[no] **cbs** *size-in-kbytes*

Context

config>qos>sap-egress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the CBS value (minimum depth of the queue).

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

Default

32

Parameters

size-in-kbytes

Specifies the CBS value, in kilobytes

Values 0 to 10240 (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
 0 to 102400 (7210 SAS-K 3SFP+ 8C)

mbs

Syntax
[no] mbs *size-in-kbytes*

Context
config>qos>sap-egress>queue

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command specifies the MBS value (maximum depth of the queue).
The **no** form of this command reverts to the default value.

Default
512

Parameters
size-in-kbytes
Specifies the MBS value, in kilobytes.

Values 0 to 12800 (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
 0 to 63488 (7210 SAS-K 3SFP+ 8C)

priority

Syntax
priority *level*
no priority

Context
config>qos>sap-egress>queue

Platforms
Supported on all 7210 SAS platforms as described in this document

Description

This command defines the queue priority. The queue priority is used by the scheduler to determine the order of service in both the within-cir loop and within-pir loop. Higher priority queues are serviced before lower priority queues.

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

Default

1

Parameters

level

Specifies the priority of the queue.

Values 1 to 4

slope-policy

Syntax

slope-policy *name*

no slope-policy

Context

config>qos>sap-egress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command is used to override the default slope policy configuration for the queue. The specified slope policy name must exist as a current slope policy name. If the slope policy does not exist, the **slope-policy** command fails. If a slope policy is currently associated with a queue, the slope policy cannot be removed from the system.

The slope policy contains the ring and non-ring high and low WRED slope definitions that are used by the queue. See [Buffer pools](#) for more information about ring and non-ring buffer pools and slope usage.

If the **slope-policy** command is not executed or the **no slope-policy** command is executed, the default slope policy will be associated with the queue.

The **no** form of this command reverts the queue to the default slope policy.

Parameters

name

Specifies an existing slope policy name, up to 32 characters.

rate

Syntax

```
rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]  
no rate
```

Context

```
config>qos>sap-egress>queue
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the administrative PIR and the administrative CIR parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets 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 max default specifies the amount of bandwidth in kilobits per second. The max value is mutually exclusive to the *pir-rate-in-kbps* value.

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 per second. The **cir** parameter overrides the default administrative CIR used by the queue.

Default	0
Values	0 to 10000000, max

pir pir-rate-in-kbps

Specifies the administrative PIR rate, in kilobits per second. When the **rate** command is executed, a PIR setting is optional.

Default	max
Values	1 to 10000000, max

weight

Syntax

weight *weight*

no weight

Context

config>qos>sap-egress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the weight of the queue.

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.

Default

1

Parameters

weight

Specifies the weight of the queue.

Values 1 to 100

scope

Syntax

scope {**exclusive** | **template**}

no scope

Context

config>qos>sap-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the scope as exclusive or template. The scope of the policy cannot be changed if the policy is applied to multiple interface 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 is 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 a template scope. An error is generated when **scope template** is changed to **scope exclusive** for default policies.

10.2.2.2 Operational commands

copy

Syntax

copy sap-egress *src-pol dst-pol* [**overwrite**]

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

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

src-pol

Specifies the destination policy.

Values 1 to 65535

overwrite

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

remark

Syntax

remark policy-id

no remark

Context

config>qos>sap-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the remarking policy ID to use for marking packets on service egress (also known as Access SAP egress).

The remarking policy ID must be associated with the appropriate SAP egress policy and remarking must be enabled in the service egress policy to enable marking of packets sent out of service (SAP) egress. Only a remarking policy of type dot1p, dscp, or dot1p-dscp is allowed when the remark policy is associated with service egress. See [Remark policies](#) for more information about remark policy types and their usage.

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 at all.

Parameters

policy-id

Specifies the remark policy.

Values 1 to 65535

remarking

Syntax

[no] remarking

Context

config>qos>sap-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

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 sap-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 SAPs.

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

The **no** form of this command disables remarking.

Default

no remarking

10.2.2.3 Show commands

sap-egress

Syntax

sap-egress [*policy-id*] [association | detail]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

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 output is an example of SAP egress QoS policy information, and [Table 45: Output fields: show SAP egress](#) describes the output fields.

Sample output

```
*A:Dut-A# show qos sap-egress 10 detail
=====
QoS Sap Egress
=====
-----
Sap Egress Policy (10)
-----
Scope                : Template
Remark               : False          Remark Pol Id       : 1
DSCP Class Policy    : 20             IP Prec Class Policy : 10
Dot1p Class Policy   : 10
Name                 : (Not Specified)
Description           : (Not Specified)
```

Table 45: Output fields: show SAP egress

Label	Description
SAP Egress Policy	Displays the ID that uniquely identifies the 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
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
DSCP Class Policy	Displays the ID that uniquely identifies the DSCP classification policy
IP Prec Class Policy	Displays the ID that uniquely identifies the DSCP classification policy with IP precedence

Label	Description
	This field applies only to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
Dot1p Class Policy	Displays the ID that uniquely identifies the dot1p classification policy
Description	A text string that helps identify the policy context in the configuration file

11 Schedulers

This chapter provides information about the scheduler support available on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.



Note:

For scheduler configuration examples, see the applicable QoS policies chapters: [Network QoS policies](#), [Network queue QoS policies](#), [Service ingress QoS policies](#), and [Service egress QoS policies](#).

11.1 Overview

Schedulers are supported on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C with strict Priority and WFQ mode of scheduling, or a mix of both, available for use.

On the 7210 SAS-K 2F1C2T, schedulers are used at SAP ingress, SAP egress, access-uplink port ingress and access-uplink port egress.


On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, schedulers are used at SAP ingress and egress; network port ingress and egress; hybrid port ingress and egress; and access-uplink port ingress and egress.

The scheduler uses two loops, the CIR loop and PIR loop, each with four priorities. The configured priority of the queue determines the service order of the queue in the CIR loop and the PIR loop. The scheduler first goes through the CIR loop, where it services all the queues which are operating at less than CIR rate according to the priority (that is, higher priority queues get services earlier than lower priority queues). It thereafter goes through the PIR loop, where it services all the queues which are operating above the CIR rate (but less than PIR rate) according to the priority (that is, higher priority queues get services earlier than lower priority queues).

If there are multiple queues configured with the same priority, in the CIR loop the queues are scheduled using WFQ, with the configured weight of the queue used to determine the proportion of the available bandwidth that is given to the queue. In the PIR loop, the queues are scheduled using WFQ, with the configured weight of the queue used to determine the proportion of the available bandwidth that is given to the queue (using WFQ). In the PIR loop, the queues are scheduled until the PIR rate is met or until no more bandwidth is available. If the PIR rate is met, then the queues are not scheduled anymore.

The following figure shows the scheduler implemented on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, or 7210 SAS-K 3SFP+ 8C.

Figure 4: Scheduler implementation



The queues at the top belong to SAP #1 and the queues at the bottom belong to SAP #2 and that all queues have the same weight. The scheduling order is (assuming unlimited bandwidth):

1. Start CIR loop.
2. Select the highest priority queues (priority of 4), that is, SAP #1 → Queue #6, SAP #2 → Queue #7 and schedule them until CIR is met.
3. Select the next highest priority queues (priority of 3), SAP #1 → Queue #5, SAP #1 → Queue #2, SAP #2 → Queue #4 and schedule them until CIR is met; Each queue is given a bandwidth in proportion to the configured weight.
4. Select the next highest priority queues (priority of 2), SAP #2 → Queue #2 and schedule it until CIR is met.
5. No more queues left to be serviced in the CIR loop; All queues CIR is met.



Note:

SAP #1 → Queue #1 and SAP #2 → Queue #1 are not serviced in the CIR loop as they are configured with CIR=0; They are serviced only in the PIR loop.

6. Start PIR loop.
7. Select the highest priority queues (priority of 4), that is, SAP #1 → Queue #6, SAP #2 → Queue #7. These do not have any PIR configured, therefore the user can skip them.
8. Select the next highest priority queues (priority of 3), SAP #1 → Queue #5, SAP #1 → Queue #2, SAP #2 → Queue #4 and schedule them until their PIR is met; Each queue is given a bandwidth in proportion to the configured weight until the PIR is met.

9. Select the next highest priority queues (priority of 2), SAP #2 → Queue #2 and schedule it until PIR is met.
10. Select the next highest priority queues (priority of 1), SAP #1 → Queue #1 and SAP #2 → Queue #1 and schedule them until the PIR is met; Each queue is given bandwidth in proportion to the configured weight until their PIR is met.

Additionally, in this section we did not take into account the port egress rate shaper or SAP aggregate shaper. These shapers limit the available bandwidth to a port or a SAP. This does not change the scheduling mechanism but provides additional controls to the user to limit the amount of bandwidth a SAP can get or a port can transmit.

On the 7210 SAS-K 2F1C2T, schedulers are available at:

- **service ingress**

This scheduler distributes the available bandwidth among all the SAPs with service ingress policies.

- **service egress**

This scheduler distributes the available bandwidth among all the SAPs configured on a specific port with service egress policies.

- **access-uplink port ingress**

This scheduler distributes the available bandwidth among up to eight ingress queues configured on the access-uplink port.

- **access-uplink port egress**

This scheduler distributes the available bandwidth among up to eight egress queues configured on the access-uplink port.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, schedulers are available at:

- **service ingress**

This scheduler distributes the available bandwidth among all the SAPs with service ingress policies.

- **service egress**

This scheduler distributes the available bandwidth among all the SAPs configured on a specific port with service egress policies.

- **network port and access-uplink port ingress**

This scheduler distributes the available bandwidth among up to eight ingress queues configured on the network/access-uplink port.

- **network port and access-uplink port egress**

This scheduler distributes the available bandwidth among up to eight egress queues configured on the network/access-uplink port.

- **hybrid port ingress**

This scheduler distributes the available bandwidth among up to eight network ingress queues configured on the hybrid port and the per-SAP ingress queues.

- **hybrid port egress**

This scheduler distributes the available bandwidth among up to eight network egress queues configured on the hybrid port and the per-SAP egress queues. The user has the option of using a SAP egress aggregate shaper rate and network queue egress aggregate shaper rate to limit the traffic per SAP and for network IP interfaces configured on the port.

12 Slope QoS policies

This chapter provides information to configure slope QoS policies using the command line interface.

12.1 Overview of buffer pools and slope policies

For an overview of buffer pools supported on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, see [Buffer pools](#).

On the 7210 SAS-K 2F1C2T, slope policies are applied to service ingress queues, service egress queues, access-uplink port ingress queues, and access-uplink port egress queues.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, slope policies are applied to service ingress and egress queues; network port ingress and egress queues; hybrid port ingress and egress queues; and access-uplink port ingress and egress queues.

Each of these queuing points allocates buffers from the buffer pool and implements WRED for congestion management. During congestion WRED is used to evaluate how buffers from the pool are allocated to different FCs and to in-profile and out-of-profile traffic within a specific FC. The slope policies define the WRED parameters to use for in-profile/high-priority packets and for out-of-profile/low-priority packets. The high-slope and low-slope define the parameters for in-profile/high-priority packets and for out-of-profile/low-priority packets respectively. In addition, on ring ports the option is available to use separate slopes for ring traffic (that is, traffic coming in on one ring port and going out of another ring port) and non-ring traffic (that is, traffic coming in on access port and going out of another access port or access-uplink port). For more information about ring and non-ring ports, see [Buffer pools](#).

12.2 Basic configurations

A basic slope QoS policy must conform to the following:

- Each slope policy must have a unique policy ID.
- High slope and low slope (both ring and non-ring slopes) are shut down by default.
- Default values can be modified but parameters cannot be deleted.

12.2.1 Creating 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 does 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.

Use the following syntax to configure a slope policy.

```
slope-policy "default" create
    description "Default slope policy."
    high-slope
        shutdown
        start-avg percent
        max-avg percent
        max-prob percent
    exit
    low-slope
        shutdown
        start-avg percent
        max-avg percent
        max-prob percent
    exit
    high-slope-ring
        shutdown
        start-avg percent
        max-avg percent
        max-prob percent
    exit
    low-slope-ring
        shutdown
        start-avg percent
        max-avg percent
        max-prob percent
    exit
```

12.2.1.1 Default slope policy values for the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Example

The following is a sample default slope policy configuration output, and [Table 46: Default slope policy values](#) lists the default values.

```
*A:SAH01-051>config>qos>slope-policy$ info detail
-----
no description
high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 80
exit
low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 80
exit
high-slope-ring
    shutdown
    start-avg 70
    max-avg 90
    max-prob 80
exit
```

```
low-slope-ring
shutdown
start-avg 50
max-avg 75
max-prob 80
exit
-----
*A:SAH01-051>config>qos>slope-policy$
```

Table 46: Default slope policy values

Description	Default slope policy
high-slope	
Administrative state	shutdown
start-avg	70% utilization
max-avg	90% utilization
max-prob	80%
low slope	
Administrative state	shutdown
start-threshold	50% utilization
max-avg	75% utilization
max-prob	80%
high-slope-ring	
Administrative state	shutdown
start-threshold	70% utilization
max-avg	90% utilization
max-prob	80%
low-slope-ring	
Administrative state	shutdown
start-threshold	50% utilization
max-avg	75% utilization
max-prob	80%

12.3 Applying slope policies

On the 7210 SAS-K 2F1C2T, slope policies are associated with service ingress queues, service egress queues, access-uplink port egress queues, and access-uplink port ingress queues.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, slope policies are associated with service ingress and egress queues; network port ingress and egress queues; hybrid port ingress and egress queues; and access-uplink port ingress and egress queues.

Example

The following syntax examples may be used to apply slope policies to service ingress queues, service egress queues, access-uplink port ingress queues, and access-uplink port egress queues on the 7210 SAS-K 2F1C2T.

```
config>qos>sap-ingress> queue id slope-policy name
config>qos>sap-egress> queue id slope-policy name
config>qos>network>ingress> queue id slope-policy name
config>qos>network-queue> queue id slope-policy name
```

Example

The following syntax examples may be used to apply slope policies to service ingress and egress queues; network port ingress and egress queues; hybrid port ingress and egress queues; and access-uplink port ingress and egress queues on the 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C.

```
config>qos>sap-ingress> queue id slope-policy name
config>qos>sap-egress> queue id slope-policy name
config>qos>network>ingress> queue id slope-policy name
config>qos>network-queue> queue id slope-policy name
```

12.4 Deleting QoS policies

A slope policy is associated by default with access and access uplink 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*. 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.

Example

```
ALA-7>config>qos# no slope-policy slopePolicy1
MINOR: QOS #1902 Slope policy has references
ALA-7>config>qos#
```

Example: Removing slope policies from ports

```
config>qos>sap-ingress> queue id no slope-policy name
config>qos>sap-egress> queue id no slope-policy name
config>qos>network>ingress> queue id no slope-policy name
config>qos>network-queue> queue id no slope-policy name
```

12.4.1 Removing a policy from the QoS configuration

Use the following to delete a slope policy.

```
config>qos# no slope-policy policy-id
```

Example

```
config>qos# no slope-policy slopePolicy1
```

12.5 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]
```

Example

The following is a sample output of copied policies.

```
*A:SAH01-051>config>qos>slope-policy# info detail
-----
      description "Default slope policy."
      high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 80
      exit
      low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 80
      exit
      high-slope-ring
        shutdown
        start-avg 70
        max-avg 90
        max-prob 80
      exit
      low-slope-ring
        shutdown
        start-avg 50
        max-avg 75
        max-prob 80
      exit
-----
*A:SAH01-051>config>qos>slope-policy#
```

12.6 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.

12.7 Slope QoS policy command reference

- [Command hierarchies](#)
- [Command descriptions](#)

12.7.1 Command hierarchies

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

12.7.1.1 Configuration commands

```
- config
  - qos
    - slope-policy name [create]
    - no slope-policy name
      - description description-string
      - no description
      - [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
      - [no] high-slope-ring
        - max-avg percent
        - no max-avg
        - max-prob percent
        - no max-prob
        - [no] shutdown
        - start-avg percent
        - no start-avg
      - [no] low-slope-ring
        - max-avg percent
```

```
- no max-avg
- max-prob percent
- no max-prob
- [no] shutdown
- start-avg percent
- no start-avg
```

12.7.1.2 Operational commands

```
- config
  - qos
    - copy slope-policy src-name dst-name [overwrite]
```

12.7.1.3 Show commands

```
- show
  - qos
    - slope-policy [slope-policy-name] [detail]
```

12.7.2 Command descriptions

12.7.2.1 Configuration commands

12.7.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>slope-policy

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

12.7.2.1.2 Slope policy QoS commands

slope-policy

Syntax

slope-policy *name* [create]

no slope-policy *name*

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

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 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.

12.7.2.1.3 Slope policy QoS policy commands

queue

Syntax

queue *queue-id*

Context

config>qos>slope-policy

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the high priority Random Early Detection (RED) slope graph. A high priority RED slope is used managing access to the shared portion of the buffer pool for high priority or in-profile packets that ingress of non-ring ports and egress out of non-ring ports.

The **high-slope** parameters can be changed at any time and the affected buffer pool high priority RED slopes are adjusted.

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 mode does 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>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the low priority RED slope graph. A low priority ring RED slope is used managing access to the shared portion of the buffer pool for low priority or out-of-profile packets that are received on non-ring ports and egress out of non-ring ports.

The **low-slope** parameters can be changed at any time and the affected buffer pool high priority RED slopes are adjusted.

The **no** form of this command reverts the low slope configuration commands to the default values. If the commands within **low-slope** are set to the default parameters, the **low-slope** mode does not appear in **save config** and **show config** output unless the detail parameter is present.

high-slope-ring

Syntax

[no] high-slope-ring

Context

config>qos>slope-policy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the high priority RED slope graph. A high priority ring RED slope is used managing access to the shared portion of the buffer pool for high priority or in-profile packets that ingress of ring ports and egress out of ring ports.

The **high-slope-ring** parameters can be changed at any time and the affected buffer pool high priority RED slopes are adjusted.

The **no** form of this command reverts the high slope configuration commands to the default values. If the commands within high-slope-ring are set to the default parameters, the high-slope node does not appear in **save config** and **show config** output unless the detail parameter is present.

low-slope-ring

Syntax

[no] low-slope-ring

Context

config>qos>slope-policy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the low priority Random Early Detection (RED) slope graph. A low priority ring RED slope is used managing access to the shared portion of the buffer pool for low priority or out-of-profile packets that are received on ring ports and egress out of ring ports.

The **low-slope-ring** parameters can be changed at any time and the affected buffer pool high priority RED slopes are adjusted.

The **no** form of this command reverts the high slope configuration commands to the default values. If the commands within low-slope-ring are set to the default parameters, the low-slope node does not appear in **save config** and **show config** output unless the detail parameter is present.

12.7.2.1.4 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

config>qos>slope-policy>queue>high-slope-ring

config>qos>slope-policy>queue>low-slope-ring

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the low priority or high priority 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 setting. If the current startavg setting is larger than the default, an error occurs and the max-avg value is not 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

max-avg 75 - Non-tcp 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 occurs and no change takes 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

config>qos>slope-policy>queue>high-slope-ring

config>qos>slope-policy>queue>low-slope-ring

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the low priority or high priority Random Early Detection (RED) slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one. The percent parameter is expressed as a percentage of packet discard probability where always discard is a probability of 1. A **max-prob** value of 75 represents 75% of 1, or a packet discard probability of 0.75.

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

Default

max-prob 75

Parameters

percent

Specifies the maximum drop probability percentage corresponding to the max-avg, expressed as a decimal integer.

Values 0 to 10, 25, 50, 75, 100

shutdown

Syntax

[no] shutdown

Context

config>qos>slope-policy>high-slope

config>qos>slope-policy>low-slope

config>qos>slope-policy>queue>high-slope-ring

config>qos>slope-policy>queue>low-slope-ring

Platforms

Supported on all 7210 SAS platforms as described in this document

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

```
config>qos>slope-policy>queue>high-slope-ring  
config>qos>slope-policy>queue>low-slope-ring
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the low priority or high priority Random Early Detection (RED) slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero. The percent parameter is expressed as a percentage of the shared buffer size.

The **no** form of this command reverts the value to the default. If the max-avg setting is smaller than the default, an error occurs and the start-avg setting is not 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

max-avg 50 - Non-tcp slope default is 50% buffer utilization

Parameters

percent

Specifies the percentage of the 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 occurs and no change takes place.

Values 0 to 100

12.7.2.1.5 Operational commands

```
copy
```

Syntax

```
copy slope-policy src-name dst-name [overwrite]
```

Context

```
config>qos
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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 src-name dst-name

Specifies the source policy ID that the **copy** command attempts to copy from and specifies the destination policy ID to which the command copies a duplicate of the policy. This parameter indicates that the source policy ID and the destination policy ID are slope policy IDs.

overwrite

Keyword to replace the existing destination policy. Everything in the existing destination policy is overwritten with the contents of the source policy. If **overwrite** is not specified, an error occurs if the destination policy ID exists.

12.7.2.1.6 Show commands

slope-policy

Syntax

slope-policy [*slope-policy-name*] [**detail**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

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 47: Output fields: QoS slope policy](#) describes the output fields.

Sample output

```
*A:SAH01-051>show>qos# slope-policy "default"
```



```
=====
QoS Slope Policy
=====
Policy      : default
Description  : Default slope policy.

-----
High Slope Parameters
-----
Start Avg   : 70                      Admin State : Disabled
Max Avg     : 90                      Max Prob.   : 80

-----
Low Slope Parameters
-----
Start Avg   : 50                      Admin State : Disabled
Max Avg     : 75                      Max Prob.   : 80

-----
High Slope Ring Parameters
-----
Start Avg   : 70                      Admin State : Disabled
Max Avg     : 90                      Max Prob.   : 80

-----
Low Slope Ring Parameters
-----
Start Avg   : 50                      Admin State : Disabled
Max Avg     : 75                      Max Prob.   : 80

=====
*A:SAH01-051>show>qos#
```

```
*A:SAH01-051>show>qos# slope-policy detail
=====
QoS Slope Policy
=====
Policy      : default
Description  : Default slope policy.

-----
High Slope Parameters
-----
Start Avg   : 70                      Admin State : Disabled
Max Avg     : 90                      Max Prob.   : 80

-----
Low Slope Parameters
-----
Start Avg   : 50                      Admin State : Disabled
Max Avg     : 75                      Max Prob.   : 80

-----
High Slope Ring Parameters
-----
Start Avg   : 70                      Admin State : Disabled
Max Avg     : 90                      Max Prob.   : 80

-----
Low Slope Ring Parameters
-----
Start Avg   : 50                      Admin State : Disabled
Max Avg     : 75                      Max Prob.   : 80

-----
SAP Ingress
-----
```

```

SAP Ingress Policy Id      : 1
Queue Ids                  : 1
-----
SAP Egress
-----
SAP Egress Policy Id      : 1
Queue Ids                  : 1
-----
Network Ingress
-----
Network Ingress Policy Id  : 1
Queue Ids                  : 1, 2, 3, 4, 5, 6, 7, 8
-----
Network Queues
-----
Network Queue Policy Name  : default
Queue Ids                  : 1, 2, 3, 4, 5, 6, 7, 8
=====
*A:SAH01-051>show>qos#

```

Table 47: 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
Admin State	Up — The administrative status of the RED slope is enabled Down — The administrative status of the RED slope is disabled Specifies the low priority or high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one
Max Prob.	Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one

13 Remark policies

This chapter provides information to configure remark policies using the CLI.

13.1 Overview of remark policies for 7210 SAS-K 2F1C2T

The remark policies configure the marking behavior for the system at the egress of access SAP or access uplink port. These policies allow the user to define the FC-to-egress marking values and to use the available hardware resources efficiently. Based on the packet encapsulation used, the remark policy allows the user to define and associate service egress policies to service egress and network QoS policies to access-uplink port. The following remark policies are available:

- **dot1p**

This policy is used for service egress policy, and network QoS.

- **dscp**

This policy is used for service egress and network QoS.

- **dot1p-dscp**

This policy is used for service egress and network QoS.

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 table.

Table 48: Summary of remark policy and attachment points for 7210 SAS-K 2F1C2T

Remark policy type	QoS policy	Attachment point	Packet header bits marked
dot1p	Service egress policy	Access SAP egress	dot1p bits and optionally the DEI bits, in the L2 header for service packets sent out of an Access SAP
	Network policy	Access-uplink port	dot1p bits and optionally the DEI bits, in the L2 header for service packets sent out of an access-uplink port (all packets sent out of all access-uplink SAPs are marked)
dscp	Service egress policy	Access SAP egress	IP DSCP bits in the IP header (if present) for service packets sent out of an Access SAP
	Network policy	Access-uplink port	IP DSCP bits in the IP header (if present) for service packets sent out of an access-uplink port Marking is done for packets sent out of all the access-uplink SAPs configured on the access-uplink port

Remark policy type	QoS policy	Attachment point	Packet header bits marked
dot1p-dscp	Service egress policy	Access SAP egress	dot1p bits and optionally the DEI bits, in the L2 header for service packets sent out of access SAP. In addition to the dot1p bits, the IP DSCP bits in the IP header (if present) are marked for service packets sent out of an Access SAP
	Network policy	Access-uplink port	dot1p bits and optionally the DEI bits, in the L2 header for service packets sent out of access uplink port. In addition to the dot1p bits, the IP DSCP bits in the IP header (if present) are marked for service packets sent out of an access-uplink port. Marking is done for packets sent out of all the access-uplink SAPs configured on the access-uplink port

13.1.1 Configuration guidelines for 7210 SAS-K 2F1C2T

The remark policy configuration guidelines for the 7210 SAS-K 2F1C2T are the following.

- The 7210 SAS-K 2F1C2T marks the dot1p configured in the marking policy, only if the node adds one or two tags or if it replaces existing tags. It does not touch the VLAN tag which represents customer payload (for example, the 7210 SAS-K 2F1C2T does not implement marking for packets when forwarding traffic out of dot1q range SAP or a null SAP).
- A limited number of unique remark policies are available to be shared among all the different types of remark policies. In other words, dot1p, dscp, and dot1p-dscp remark policy types share the available resources and scaling one type of remark policy reduces the amount policies allowed for other type of remark policy.
- The user also has an option to preserve the dot1p values in the received packet. See the CLI description for the **config service sap egress dot1p-inner** and **config service sap egress dot1p-outer** CLI commands in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Services Guide*.

13.2 Overview of remark policies for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The remark policies are used to configure the marking behavior for the system at the egress of access SAPs, network port, hybrid ports, and access uplink ports. These policies allow the user to define the FC-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 service egress policies to access SAP egress and network QoS policies applied to network ports, hybrid ports, and access-uplink ports. The following type of remark policies are available:

- **dot1p**
This policy is used for service egress and network QoS policies.
- **dscp**

This policy is used for service egress and network QoS policies.

- **dot1p-dscp**

This policy is used for service egress and network QoS policies.

- **lsp-exp**

This policy is used for network QoS policies (only network ports and hybrid ports).

- **dot1p-lsp-exp**

This policy is used for network QoS policies (only network ports and hybrid ports).

- **dot1p-lsp-exp-dscp**

This policy is used for network QoS policies (only network ports and hybrid ports).

- **dscp-lsp-exp**

This policy is used for network QoS policies (only network ports and hybrid ports).

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 [Table 49: Summary of remark policy and attachment points for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#).

Table 49: Summary of remark policy and attachment points for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Remark policy type	QoS policy	Attachment point	Packet header bits marked
dot1p	Service egress policy	Access SAP egress (applicable to SAPs configured on access ports or hybrid ports)	dot1p bits, and optionally the DEI bits, in the outermost Ethernet header for service packets sent out of an Access SAP
	Network policy	Network port hybrid port	dot1p bits, and optionally the DEI bits, in the outermost Ethernet header for all MPLS and IP packets sent out of a network IP interface configured on a network port or hybrid port
	Network policy	Access-uplink port	dot1p bits, and optionally the DEI bits, in the Ethernet header for service packets sent out of an access-uplink port (all packets sent out of all access-uplink SAPs are marked)
dscp	Service egress policy	Access SAP egress (applicable to SAPs configured on access ports or hybrid ports)	IP DSCP bits in the IP header (if present) for service packets sent out of an Access SAP
	Network policy	Network port hybrid port	IP DSCP bits in the IP header (if present) for all IP packets sent out of a network IP interface configured on a network port or hybrid port

Remark policy type	QoS policy	Attachment point	Packet header bits marked
		Access-uplink port	IP DSCP bits in the IP header (if present) for service packets sent out of an Access-uplink port. Marking is done for packets sent out of all access-uplink SAPs configured on the access-uplink port
dot1p-dscp	Service egress policy	Access SAP egress (applicable to SAPs configured on access ports or hybrid ports)	dot1p bits, and optionally the DEI bits, in the L2 header for service packets sent out of access SAP. In addition to the dot1p bits, the IP DSCP bits in the IP header (if present) are marked for service packets sent out of an Access SAP
	Network policy	Network port hybrid port	dot1p bits, and optionally the DEI bits, in the outermost Ethernet header for all MPLS and IP packets sent out of a network IP interface configured on a network port or hybrid port IP DSCP bits in the IP header (if present) for all IP packets sent out of a network port IP interface configured on a network port or hybrid port
		Access-uplink port	dot1p bits, and optionally the DEI bits, in the Ethernet header for service packets sent out of access uplink port. In addition to the dot1p bits, the IP DSCP bits in the IP header (if present) are marked for service packets sent out of an access-uplink port. Marking is done for packets sent out of all the access-uplink SAPs configured on the access-uplink port
lsp-exp	Network policy	Network port hybrid port	MPLS EXP bits for the outermost MPLS labels for all MPLS packets sent out of a network IP interface configured on a network port or hybrid port
dot1p-lsp-exp	Network policy	Network port hybrid port	dot1p bits, and optionally the DEI bits, in the outermost Ethernet header for all MPLS and IP packets sent out of a network port MPLS EXP bits for the outermost MPLS labels for all MPLS packets sent out of a network IP interface configured on a network port or hybrid port
dot1p-lsp-exp-dscp	Network policy	Network port hybrid port	dot1p bits, and optionally the DEI bits, in the outermost Ethernet header for all MPLS and IP packets sent out of a network IP interface configured on a network port or hybrid port

Remark policy type	QoS policy	Attachment point	Packet header bits marked
			IP DSCP for all IP packets sent out of a network IP interface configured on a network port or hybrid port MPLS EXP bits for the outermost MPLS labels for all MPLS packets sent out of a network IP interface configured on a network port or hybrid port
dscp-lsp-exp	Network Policy	Network port hybrid port	IP DSCP for all IP packets sent out of a network IP interface configured on a network port or hybrid port MPLS EXP bits for the outermost MPLS labels for all MPLS packets sent out of a network IP interface configured on a network port or hybrid port

13.2.1 Configuration guidelines for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The remark policy configuration guidelines for the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C are the following:

- The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C mark the dot1p configured in the marking policy only if the node adds one or two tags or if it replaces existing tags. The platforms do not modify the VLAN tag which represents customer payload; for example, the platforms do not mark the packets when forwarding traffic out of dot1q range SAP or a null SAP.
- A limited number of unique remark policies are available to be shared among all the different types of remark policies. In other words, dot1p, dscp, dot1p-dscp, lsp-exp, dot1p-lsp-exp, dot1p-dscp-lsp-exp remark policy types share the available resources and scaling one type of remark policy reduces the amount policies allowed for other type of remark policy.
- The user also has an option to preserve the dot1p values in the received packet. See the CLI description for the **config service sap egress dot1p-inner** and **config service sap egress dot1p-outer** commands in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Services Guide*.
- When marking is enabled without explicitly associating the remark policy, the default remark policy is used for marking packets.
- Any remark policy that specifies MPLS lsp-exp values cannot be associated with access-uplink port. The association is blocked by the system.
- MPLS packets sent out of a network IP interface configured on a network port or hybrid port can be marked with MPLS lsp-exp values and dot1p, DEI values if both are configured.
- IP packets sent out of a network IP interface configured on a network port or hybrid port can be marked with IP DSCP values and dot1p, DEI values if both are configured.
- The dot1p bits of the packets sent out of a network IP interface configured on a network port or hybrid port are always marked by default, even when remarking is disabled.

13.3 Basic configurations

A basic remark policy must conform to the following:

- Each remark policy must have a unique policy ID.
- The remark policy type must be specified.
- The FC-to-egress marking values must be specified.

13.3.1 Creating a remark policy

To create a new remark policy, define the following:

- a remark policy name and type (this is optional and by default is dot1p)
- a brief description of the policy features
- the FC-to-egress marking values

Use the following syntax to configure a remark policy.

```
-----
*A:7210SAS>config>qos>remark# info detail (applicable to SAS-K devices)
-----
no description
fc af
    de-mark-outer
    dotlp-inner 4
    dotlp-outer in-profile 2 out-profile 3
exit
fc be
    de-mark-outer force 0
    dotlp-inner in-profile 4 out-profile 2
    dotlp-outer 4
exit
fc ef
    no de-mark-outer
    dotlp-inner in-profile 4 out-profile 5
    dotlp-outer in-profile 3 out-profile 1
exit
fc h1
    de-mark-outer
    dotlp-inner 1
    dotlp-outer in-profile 1 out-profile 2
exit
fc h2
    de-mark-outer
    dotlp-inner in-profile 4 out-profile 4
    dotlp-outer in-profile 6 out-profile 3
exit
fc l1
    de-mark-outer force 1
    dotlp-inner in-profile 2 out-profile 6
    dotlp-outer 4
exit
fc l2
    de-mark-outer
    dotlp-inner 7
    dotlp-outer 3
```



```
        exit
        fc nc
            no de-mark-outer
            dot1p-inner in-profile 2 out-profile 4
            dot1p-outer in-profile 5 out-profile 6
        exit
*A:7210SAS>config>qos>remark#
```

13.3.2 Editing QoS policies

About this task

Existing policies and entries can be edited using 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.

13.4 Remark policy command reference

- [Command hierarchies](#)
- [Command descriptions](#)

13.4.1 Command hierarchies

- [Configuration commands for 7210 SAS-K 2F1C2T](#)
- [Configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [Operational commands](#)
- [Show commands](#)

13.4.1.1 Configuration commands for 7210 SAS-K 2F1C2T

```
config
- qos
  - no remark policy-id
  - remark policy-id [create] [remark-type remarking-type]
  - [no] description description-string
  - [no] fc fc-name
    - [no] de-mark-outer [force de-value]
    - dot1p-inner dot1p-value
    - dot1p-inner [in-profile dot1p-value] [out-profile dot1p-value]
    - no dot1p-inner
    - no dot1p-outer
    - dot1p-outer dot1p-value
```

```
- dot1p-outer [in-profile dot1p-value] [out-profile dot1p-value]
- dscp-in-profile dscp-name
- no dscp-in-profile
- dscp-out-profile dscp-name
- no dscp-out-profile
```

13.4.1.2 Configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
config
- qos
  - no remark policy-id
  - remark policy-id [create] [remark-type remarking-type]
  - [no] description description-string
  - [no] fc fc-name
    - [no] de-mark-outer [force de-value]
    - dot1p-inner dot1p-value
    - dot1p-inner [in-profile dot1p-value] [out-profile dot1p-value]
    - no dot1p-inner
    - no dot1p-outer
    - dot1p-outer dot1p-value
    - dot1p-outer [in-profile dot1p-value] [out-profile dot1p-value]
    - 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
    - lsp-exp-out-profile lsp-exp-value
```

13.4.1.3 Operational commands

```
config
- qos
  - copy remark src-pol dst-pol [overwrite]
```

13.4.1.4 Show commands

```
show
- qos
  - remark-policy [policy-id] [association | detail]
```

13.4.2 Command descriptions

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

13.4.2.1 Configuration commands

13.4.2.1.1 Generic commands

description

Syntax

[no] description *description-string*

Context

config>qos>remark

Platforms

Supported on all 7210 SAS platforms as described in this document

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 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.

13.4.2.1.2 Remark policy QoS commands

remark

Syntax

no remark *policy-id*

remark *policy-id* [**create**] [**remark-type** *marking-type*]

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a new remark policy of the specified type.

The following types of remark policies are available on the 7210 SAS-K 2F1C2T:

- dot1p
- dscp
- dot1p-dscp

The following type of remark policies are available on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C:

- dot1p — used for service egress and network QoS policies
- dscp — used for service egress and network QoS policies
- dot1p-dscp — used for service egress and network QoS policies
- lsp-exp — used for network QoS policies (only network ports and hybrid ports)
- dot1p-lsp-exp — used for network QoS policies (only network ports and hybrid ports)
- dot1p-lsp-exp-dscp — used for network QoS policies (only network ports and hybrid ports)
- dscp-lsp-exp — used for network QoS policies (only network ports and hybrid ports)

The *marking-type* of the policy also determines the values user is allowed to configure in the policy and the QoS policy with which this policy can be associated with. See the [Table 48: Summary of remark policy and attachment points for 7210 SAS-K 2F1C2T](#) for the different remark policies supported on the node and its use.



Note:

A limited number of unique remark policies are available to be shared among all the different types of remark policies. The remark policy types share the available resources, and scaling one type of remark policy reduces the number of policies allowed for other types of remark policies.

Default

no default

Parameters

policy-id

Specifies the policy ID of the remark policy.

Values 1 to 65535

marking-type

Specifies the type of marking values in the remark policy.

Values dot1p — Specify FC to 802.1 Dot1p value to use for marking. It is the default used if the user does not explicitly specify the marking-type value. dscp — Specify FC to IP DSCP value to use for marking.

dot1p-dscp - Specify FC to both Dot1p and IP DSCP values to use for marking.

Values lsp-exp — Available only on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C. Specify FC to MPLS EXP values to use for marking.

Values dot1p-lsp-exp — Available only on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C. Specify FC to MPLS EXP values and Dot1p values to use for marking.

Values dot1p-dscp-lsp-exp — Available only on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C. Specify FC to MPLS EXP values, IP DSCP values and Dot1p values to use for marking.

fc

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

Context
config>qos>remark

Platforms
Supported on all 7210 SAS platforms as described in this document

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 forwarding class 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

13.4.2.1.3 Remark policy forwarding class commands

de-mark-outer

Syntax

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

Context

config>qos>remark-policy>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command explicitly defines the marking of the DEI bit of the outermost tag for **fc** *fc-name* according to the in and out of profile status of the packet.

If no *de-value* is present, the default values are used for the marking of the DE bit, 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 or out of profile status.



Note:

If remarking is enabled, the inner tag DEI is always set to zero irrespective of de-mark-outer.

If remarking is enabled and de-mark-outer is not configured then, DEI bit of the outer tag is set to zero.

If remarking is disabled then, the DEI bits are preserved for both inner and outer tag.

Parameters

de-value

Specifies the DEI value to use for this forwarding class.

Values 0 or 1

dot1p-inner

Syntax

no dot1p-inner

dot1p-inner *dot1p-value*

Context

config>qos>remark-policy>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command explicitly defines the egress IEEE 802.1P (dot1p) bits for the inner VLAN tag marking for *fc-name*. When the marking is set, all packets of *fc-name* that egresses out of QinQ SAP only (that is, SAPs configured with two VLAN tags explicitly defined; for example, SAP 1/1/5:10.100) use the explicitly defined *dot1p-value*.

This command has no effect for egress packets sent out of all other non-QinQ SAPs; for example, 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.

If the **no** form of this command is executed, the default remarking values are used.

Default

no dot1p-inner

Parameters

dot1p-value

Specifies the 802.1p value to set for the 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-policy>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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 (that is, SAPs configured with two VLAN tags explicitly defined; for example, SAP 1/1/5:10.100) use the explicitly defined in-profile and out-profile *dot1p-value*.

This command adds the capability to mark the in and out profile status on an egress through a specific dot1p combination. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network or customer device that does not support the DEI bit.

This command has no effect for egress packets sent out of all other non-QinQ SAPs; for example, 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 the **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

dot1p-outer

Syntax

dot1p-outer *dot1p-value*

no dot1p-outer

Context

config>qos>remark-policy>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the **dot1p** command has no effect.

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

Default

no dot1p-outer

Parameters

dot1p-value

Specifies the dot1p values to use.

Values 0 to 7

dscp-in-profile

Syntax

dscp-in-profile *dscp-name*

no dscp-in-profile

Context

config>qos>remark-policy>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

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 overwrites the previous value.

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

Parameters

dscp-name

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, cp, cp61, cp62, cp63

dscp-out-profile

Syntax

dscp-out-profile *dscp-name*

no dscp-out-profile

Context

config>qos>remark-policy>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

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 overwrites the previous value.

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

Parameters

dscp-name

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, cp, cp61, cp62, cp63
---------------	---

dot1p-outer

Syntax

[no] dot1p-outer [in-profile *dot1p-value*] [out-profile *dot1p-value*]

Context

config>qos>remark-policy>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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 in-profile *dot1p-value*. Similarly all the packet with out-of-profile status use the explicitly defined out-profile *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, this command has no effect.

This command adds the capability to mark on an egress the in and out profile status via a specific 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 or customer device that does not support the DEI bit.

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 DEI bit is set to 0.

This variant of the command is mutually exclusive to the use of the **dot1p-outer** command. In other words, user has a choice to use either this command or the **dot1p-outer** command but not both together.

If the **no** form of this command is executed then the default remark values are used.

Default

no dot1p-outer

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

lsp-exp-in-profile

Syntax

[no] **lsp-exp-in-profile** *lsp-exp-value*

Context

config>qos>remark-policy>fc

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the in-profile MPLS EXP value for the forwarding class. When marking is set, the corresponding MPLS EXP value is used to mark all MPLS packets with in-profile status, on the egress of this forwarding class queue.

When multiple MPLS EXP names are associated with the forwarding class in the policy, the last name entered overwrites the previous value.

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

Default

no lsp-exp-in-profile

Parameters

lsp-exp-value

Specifies the in-profile MPLS EXP value.

Values 0 to 7

lsp-exp-out-profile

Syntax

[no] **lsp-exp-out-profile** *lsp-exp-value*

Context

config>qos>remark-policy>fc

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the out-of-profile MPLS EXP value for the forwarding class. When marking is set, the corresponding MPLS EXP value is used to mark all MPLS packets with out-of-profile status, on the egress of this forwarding class queue.

When multiple MPLS EXP names are associated with the forwarding class in the policy, the last name entered overwrites the previous value.

The **no** form of this command reverts to the factory default out-profile MPLS EXP value setting for policy ID 1.

Default

no lsp-exp-out-profile

Parameters

lsp-exp-value

Specifies the out-of-profile MPLS EXP value.

Values 0 to 7

13.4.2.2 Operational commands

copy

Syntax

copy remark *src-pol dst-pol* [**overwrite**]

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

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 key word overwrite 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 source policy.

13.4.2.3 Show commands

remark-policy

Syntax

remark-policy [*policy-id*] [**association** | **detail**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

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 50: Output fields: remark policy](#) describes the output fields.

Sample output

```
*A:SAH01-051> show qos remark-policy 300 detail
=====
QoS Remarking Policies
=====
-----
Remark Policy-id      : 300          Type          : dscp
Description           : (Not Specified)
-----
FC Name      DSCP      DSCP
             InProf     OutProf
-----
be           be         be
l2           cs1        cs1
af           af11       af12
l1           af21       af22
h2           af41       af41
ef           ef         ef
h1           nc1        nc1
nc           nc2        nc2
-----
Associations
-----
SAP Egress
-----
No SAP Egress Associations found.
-----
Network
-----
No Network Policy Associations found.
=====
```

```
*A:SAH01-051>show>qos#
```

Table 50: Output fields: remark policy

Label	Description
Remark Policy ID	The ID that uniquely identifies the policy
Type	Displays the type of the remark policy
dot1P	Dot1p value for in-profile packets
fc name	Forwarding class name

14 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.

14.1 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

14.2 Certificate Management

RFC 5280, Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile 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 7468, Textual Encodings of PKIX, PKCS, and CMS Structures 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



Note:

Only PKIX and PKCS is supported.

14.3 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, Sx/S-1/10GE, and Sx/S-1/10GE-VC



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, Sx/S-1/10GE, and Sx/S-1/10GE-VC



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

14.4 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-/10GE standalone mode only.

draft-ietf-bess-evpn-vpws-14, Virtual Private Wire Service support in Ethernet VPN is supported on Mxp

14.5 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.

14.6 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 public key 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 5246, The Transport Layer Security (TLS) Protocol Version 1.2 – TLS client, RSA public key 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



Note:

Only TLS client.

RFC 5425, Transport Layer Security (TLS) Transport Mapping for Syslog – RFC 3164 with TLS 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



Note:

Only TLS client.

RFC 5656, Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer – ECDSA 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 7301, Transport Layer Security (TLS) Application Layer Protocol Negotiation Extension 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



Note:

Only TLS client.

RFC 8446, The Transport Layer Security (TLS) Protocol Version 1.3 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



Note:

Only TLS client.

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

14.7 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 is 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.

14.8 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.

14.9 IP — Version 6

RFC 2464, Transmission of IPv6 Packets over Ethernet 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



Note:

D, Dxp, K5, 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, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, K5, and T(A) for Management only.

RFC 4007, IPv6 Scoped Address 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



Note:

D, Dxp, K5, and T(A) for Management only.

RFC 4193, Unique Local IPv6 Unicast 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



Note:

D, Dxp, K5, and T(A) for Management only.

RFC 4291, Internet Protocol Version 6 (IPv6) Addressing 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



Note:

D, Dxp, K5, 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, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, K5, and T(A) for Management only.

RFC 4861, Neighbor Discovery for IP version 6 (IPv6) 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:

D, Dxp, K5, and T(A) for Management only.

RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) 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:

D, Dxp, K5, and T(A) for Management only.

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, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, K5, 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

14.10 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.

14.11 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

14.12 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

ianaiftype-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

14.13 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

14.14 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

14.15 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.

14.16 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

14.17 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

14.18 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

14.19 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

14.20 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

14.21 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

14.22 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.

14.23 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, 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

14.24 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](#)