



7210 Service Access System

Release 25.9.R1

7210 SAS-D, Dxp Quality of Service Guide

3HE 21175 AAAB TQZZA 01

Edition: 01

September 2025

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

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

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

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

No part of this document may be modified.

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

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

© 2025 Nokia.

Table of contents

| | |
|---|-----------|
| List of tables..... | 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 Nokia 7210 SAS-series services 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..... | 22 |
| | |
| 2 QoS policies..... | 24 |
| 2.1 QoS policies overview..... | 24 |
| 2.2 Overview of QoS policies on 7210 SAS-D and 7210 SAS-Dxp..... | 25 |
| 2.2.1 Summary of major functions of QoS policies..... | 26 |
| 2.3 Network and service QoS policies..... | 27 |
| 2.3.1 Network QoS policies on access-uplink ports..... | 27 |
| 2.3.2 Default network QoS policy (egress) for the 7210 SAS-Dxp..... | 29 |
| 2.3.3 Default network QoS policy (ingress)..... | 30 |
| 2.3.4 Network queue policies in access-uplink mode..... | 30 |
| 2.3.5 Service ingress QoS policies..... | 32 |
| 2.3.5.1 Default service ingress policy..... | 33 |
| 2.3.6 Service ingress classification..... | 34 |
| 2.3.6.1 Service ingress classification..... | 34 |
| 2.3.6.2 Hierarchical ingress policing on 7210 SAS-D and 7210 SAS-Dxp..... | 35 |
| 2.4 Access egress QoS policies..... | 35 |
| 2.4.1 Default access egress policy..... | 36 |
| 2.5 Meters/policers..... | 37 |
| 2.5.1 Metering/policing and meter parameters..... | 37 |
| 2.5.1.1 Meter ID..... | 37 |
| 2.5.1.2 Committed information rate (meters)..... | 37 |

| | | |
|----------|---|----|
| 2.5.1.3 | Peak information rate (meters)..... | 37 |
| 2.5.1.4 | Adaptation rule for meters..... | 38 |
| 2.5.1.5 | Committed burst size (meters/policers)..... | 39 |
| 2.5.1.6 | Maximum burst size (meters/policers)..... | 40 |
| 2.5.2 | Meter counters..... | 40 |
| 2.5.3 | Meter modes..... | 41 |
| 2.5.3.1 | Color-aware and color-blind policing/metering..... | 41 |
| 2.5.4 | QoS overrides for meters/policers..... | 42 |
| 2.5.4.1 | Configuration guidelines for QoS override..... | 42 |
| 2.5.4.2 | Configuring meter override parameters..... | 42 |
| 2.6 | Forwarding classes..... | 43 |
| 2.6.1 | FC-to-queue ID map..... | 44 |
| 2.7 | Schedulers..... | 44 |
| 2.7.1 | Port scheduler policies..... | 44 |
| 2.7.2 | Scheduler modes..... | 45 |
| 2.8 | CPU queues..... | 46 |
| 2.8.1 | Overview..... | 46 |
| 2.9 | Egress port rate limiting..... | 47 |
| 2.10 | Queue management..... | 47 |
| 2.10.1 | Queues and queue parameters..... | 47 |
| 2.10.1.1 | Queue ID..... | 47 |
| 2.10.1.2 | Committed information rate..... | 47 |
| 2.10.1.3 | Peak information rate..... | 48 |
| 2.10.1.4 | Adaptation rule for queues..... | 49 |
| 2.10.1.5 | Committed burst size (queue)..... | 50 |
| 2.10.1.6 | Maximum burst size (queue)..... | 50 |
| 2.10.1.7 | Queue counters..... | 52 |
| 2.10.2 | Buffer pools..... | 52 |
| 2.10.2.1 | Buffer pools..... | 52 |
| 2.10.2.2 | Decommissioning ports with per-port MBS pool on 7210 SAS-Dxp..... | 52 |
| 2.10.3 | Slope policies..... | 54 |
| 2.10.4 | RED slopes..... | 54 |
| 2.10.4.1 | Operation and configuration of RED slopes for 7210 SAS-D..... | 54 |
| 2.10.4.2 | Operation and configuration of RED slopes for 7210 SAS-Dxp..... | 55 |
| 2.10.4.3 | Simplified overview of RED for 7210 SAS-D and 7210 SAS-Dxp..... | 55 |

| | | |
|----------|--|-----------|
| 2.10.4.4 | Tuning the shared buffer utilization calculation on 7210 SAS-D and 7210 SAS-Dxp..... | 57 |
| 2.10.4.5 | Slope policy parameters for 7210 SAS-D..... | 58 |
| 2.10.4.6 | Slope policy parameters for 7210 SAS-Dxp..... | 59 |
| 2.11 | QoS policy entities..... | 60 |
| 2.12 | Configuration notes..... | 61 |
| 3 | DEI-based classification and marking..... | 62 |
| 3.1 | DEI-based classification..... | 62 |
| 3.2 | DEI-based marking..... | 63 |
| 3.3 | Configuration guidelines..... | 63 |
| 4 | Port level egress rate-limiting..... | 64 |
| 4.1 | Overview..... | 64 |
| 4.1.1 | Applications..... | 64 |
| 4.1.2 | Effect of port-level rate-limiting on queue functionality..... | 64 |
| 4.2 | Basic configurations..... | 65 |
| 4.2.1 | Modifying port level egress-rate command..... | 65 |
| 4.2.2 | Removing port level egress-rate command..... | 66 |
| 4.2.2.1 | Default egress-rate values..... | 66 |
| 4.3 | Port level egress-rate command reference..... | 66 |
| 4.3.1 | Command hierarchies..... | 66 |
| 4.3.1.1 | Configuration commands..... | 66 |
| 4.3.1.2 | Show commands..... | 66 |
| 4.3.2 | Configuration descriptions..... | 67 |
| 4.3.2.1 | Configuration commands..... | 67 |
| 4.3.2.2 | Show commands..... | 68 |
| 5 | SAP egress aggregate meter..... | 75 |
| 5.1 | Overview..... | 75 |
| 5.1.1 | Configuration notes..... | 75 |
| 5.2 | Basic configurations..... | 75 |
| 5.3 | SAP egress aggregate meter command reference..... | 76 |
| 6 | Frame-based accounting..... | 77 |
| 6.1 | Overview..... | 77 |
| 6.1.1 | Frame-based accounting..... | 77 |

| | | |
|----------|---|-----------|
| 6.1.2 | Effects of enabling ingress frame-based accounting on ingress meter functionality..... | 77 |
| 6.1.3 | Effects of enabling egress frame-based accounting on access-uplink queue functionality..... | 77 |
| 6.1.4 | Accounting and statistics..... | 78 |
| 6.2 | Basic configurations..... | 78 |
| 6.2.1 | Enabling and disabling frame-based accounting..... | 78 |
| 6.3 | Frame-based accounting command reference..... | 79 |
| 6.3.1 | Command hierarchies..... | 79 |
| 6.3.1.1 | Configuration commands..... | 79 |
| 6.3.1.2 | Show commands..... | 79 |
| 6.3.2 | Configuration descriptions..... | 79 |
| 6.3.2.1 | Configuration commands..... | 80 |
| 6.3.2.2 | Show commands..... | 81 |
| 7 | Network QoS policies..... | 91 |
| 7.1 | Overview of network QoS policies..... | 91 |
| 7.1.1 | Resource allocation for network QoS policy..... | 92 |
| 7.1.1.1 | Network QoS policies resource usage examples..... | 93 |
| 7.1.2 | Basic configuration..... | 100 |
| 7.1.3 | Create a network QoS policy..... | 100 |
| 7.1.4 | Default network policy values..... | 102 |
| 7.1.5 | DSCP marking for CPU-generated traffic..... | 102 |
| 7.1.6 | Default DSCP mapping table..... | 103 |
| 7.2 | Service management tasks..... | 104 |
| 7.2.1 | Deleting QoS policies..... | 104 |
| 7.2.2 | Remove a policy from the QoS configuration..... | 105 |
| 7.2.3 | Copying and overwriting network policies..... | 105 |
| 7.2.4 | Editing QoS policies..... | 105 |
| 7.3 | Network QoS policy command reference..... | 106 |
| 7.3.1 | Command hierarchies..... | 106 |
| 7.3.1.1 | Configuration commands for 7210 SAS-D and 7210 SAS-Dxp..... | 106 |
| 7.3.1.2 | Operational commands..... | 107 |
| 7.3.1.3 | Show commands..... | 107 |
| 7.4 | Command descriptions..... | 107 |
| 7.4.1 | Configuration commands..... | 107 |
| 7.4.1.1 | Generic commands..... | 107 |

| | | |
|----------|---|------------|
| 7.4.1.2 | Operational commands..... | 108 |
| 7.4.1.3 | Network QoS policy commands..... | 110 |
| 7.4.1.4 | Network ingress QoS policy commands..... | 111 |
| 7.4.1.5 | Network egress QoS policy commands..... | 123 |
| 7.4.1.6 | Network egress QoS policy forwarding class commands..... | 128 |
| 7.4.1.7 | Show commands..... | 131 |
| 8 | Network queue QoS policies..... | 137 |
| 8.1 | Overview..... | 137 |
| 8.2 | Basic configurations..... | 137 |
| 8.2.1 | Creating a network queue QoS policy..... | 137 |
| 8.2.2 | Applying network queue configuration in access-uplink mode..... | 139 |
| 8.3 | Default network queue policy values..... | 139 |
| 8.4 | Service management tasks..... | 143 |
| 8.4.1 | Deleting network queue QoS policies..... | 143 |
| 8.4.2 | Copying and overwriting network queue QoS policies..... | 144 |
| 8.4.3 | Editing network queue QoS policies..... | 145 |
| 8.5 | Network queue QoS policy command reference..... | 145 |
| 8.5.1 | Command hierarchies..... | 145 |
| 8.5.1.1 | Configuration commands..... | 145 |
| 8.5.1.2 | Operational commands..... | 146 |
| 8.5.1.3 | Show commands..... | 146 |
| 8.6 | Command descriptions..... | 146 |
| 8.6.1 | Configuration commands..... | 146 |
| 8.6.1.1 | Generic commands..... | 146 |
| 8.6.1.2 | Operational commands..... | 147 |
| 8.6.1.3 | Network queue QoS policy commands..... | 147 |
| 8.6.1.4 | Network queue QoS policy queue commands..... | 148 |
| 8.6.1.5 | Show commands..... | 152 |
| 9 | Service ingress QoS policies..... | 155 |
| 9.1 | Overview of service ingress policy..... | 155 |
| 9.1.1 | Default SAP ingress policy..... | 155 |
| 9.1.1.1 | SAP ingress policy defaults..... | 155 |
| 9.1.1.2 | Use of index file by SAP QoS ingress policy..... | 156 |
| 9.1.1.3 | Service ingress meter selection rules..... | 158 |

| | | |
|-----------|--|------------|
| 9.1.1.4 | Service ingress QoS policy configuration considerations..... | 160 |
| 9.1.2 | Resource allocation for service ingress QoS policy classification rules..... | 161 |
| 9.1.3 | Computation of SAP ingress classification and meter resources used per SAP ingress policy..... | 163 |
| 9.1.3.1 | Determining the number of classification entries..... | 164 |
| 9.1.3.2 | Service ingress QoS policies resource usage examples..... | 166 |
| 9.2 | Basic configurations..... | 187 |
| 9.2.1 | Create service ingress QoS policies..... | 187 |
| 9.2.1.1 | Service ingress QoS policy..... | 187 |
| 9.2.1.2 | Applying service ingress policies..... | 191 |
| 9.3 | Service management tasks..... | 193 |
| 9.3.1 | Deleting QoS policies..... | 193 |
| 9.3.1.1 | Removing a QoS policy from service SAPs..... | 193 |
| 9.3.2 | Copying and overwriting QoS policies..... | 194 |
| 9.3.3 | Remove a policy from the QoS configuration..... | 194 |
| 9.3.4 | Editing QoS policies..... | 194 |
| 9.4 | Service SAP QoS policy command reference..... | 194 |
| 9.4.1 | Service ingress QoS policy commands..... | 194 |
| 9.4.2 | Operational commands..... | 196 |
| 9.4.3 | Show commands..... | 196 |
| 9.5 | Command descriptions..... | 196 |
| 9.5.1 | Configuration commands..... | 196 |
| 9.5.1.1 | Generic commands..... | 196 |
| 9.5.1.2 | Operational commands..... | 197 |
| 9.5.1.3 | Service ingress QoS policy commands..... | 199 |
| 9.5.1.4 | Service ingress QoS policy forwarding class commands..... | 207 |
| 9.5.1.5 | Service ingress QoS policy entry commands..... | 210 |
| 9.5.1.6 | IP QoS policy match commands..... | 214 |
| 9.5.1.7 | Service ingress MAC QoS policy match commands..... | 220 |
| 9.5.1.8 | Service meter QoS policy commands..... | 224 |
| 9.5.1.9 | Show commands..... | 232 |
| 10 | Access egress QoS policies..... | 239 |
| 10.1 | Overview..... | 239 |
| 10.1.1 | Basic configurations..... | 239 |
| 10.1.1.1 | Modifying access egress QoS queues..... | 239 |

| | | |
|-----------|---|------------|
| 10.1.1.2 | Applying access egress QoS policies..... | 240 |
| 10.1.1.3 | Default access egress QoS policy values..... | 241 |
| 10.1.1.4 | Deleting QoS policies..... | 242 |
| 10.1.1.5 | Removing a policy from the QoS configuration..... | 242 |
| 10.2 | Access egress QoS policy command reference..... | 243 |
| 10.2.1 | Command hierarchies..... | 243 |
| 10.2.1.1 | Configuration commands for 7210 SAS-D and 7210 SAS-Dxp..... | 243 |
| 10.2.1.2 | Show commands..... | 243 |
| 10.3 | Command descriptions..... | 244 |
| 10.3.1 | Configuration commands..... | 244 |
| 10.3.1.1 | Generic commands..... | 244 |
| 10.3.1.2 | Access egress queue QoS policy commands..... | 253 |
| 11 | Port scheduler policies..... | 260 |
| 11.1 | Configuring port scheduler policies..... | 260 |
| 11.2 | Basic configurations..... | 260 |
| 11.2.1 | Creating a QoS port scheduler policy..... | 260 |
| 11.3 | Service management tasks..... | 261 |
| 11.3.1 | Copying and overwriting port scheduler policies..... | 261 |
| 11.3.2 | Editing QoS policies..... | 262 |
| 11.4 | QoS port scheduler policy command reference..... | 262 |
| 11.4.1 | Command hierarchies..... | 262 |
| 11.4.1.1 | Port scheduler policy configuration commands..... | 262 |
| 11.4.1.2 | Operational commands..... | 263 |
| 11.4.1.3 | Show commands..... | 263 |
| 11.5 | Command descriptions..... | 263 |
| 11.5.1 | Configuration commands..... | 263 |
| 11.5.1.1 | Generic commands..... | 263 |
| 11.5.1.2 | Operational commands..... | 264 |
| 11.5.1.3 | Port scheduler policy commands..... | 264 |
| 11.5.1.4 | Show commands..... | 267 |
| 12 | Slope QoS policies..... | 270 |
| 12.1 | Overview of buffer pools..... | 270 |
| 12.1.1 | Configuration guidelines for 7210 SAS-D..... | 270 |
| 12.1.2 | Configuration guidelines for 7210 SAS-Dxp..... | 271 |

| | | |
|-----------|---|------------|
| 12.2 | Basic configurations..... | 271 |
| 12.2.1 | Create a slope QoS policy for 7210 SAS-D..... | 272 |
| 12.2.2 | Create a slope QoS policy for 7210 SAS-Dxp..... | 273 |
| 12.3 | Applying slope policies..... | 274 |
| 12.3.1 | Default slope policy values..... | 274 |
| 12.3.1.1 | Default slope values for 7210 SAS-D..... | 274 |
| 12.3.1.2 | Default slope values for 7210 SAS-Dxp..... | 277 |
| 12.4 | Deleting QoS policies..... | 280 |
| 12.4.1 | Remove a policy from the QoS configuration..... | 281 |
| 12.5 | Copying and overwriting QoS policies..... | 281 |
| 12.6 | Editing QoS policies..... | 286 |
| 12.7 | Slope QoS policy command reference..... | 286 |
| 12.7.1 | Command hierarchies..... | 286 |
| 12.7.1.1 | Configuration commands for 7210 SAS-D..... | 286 |
| 12.7.1.2 | Configuration commands for 7210 SAS-Dxp..... | 287 |
| 12.7.1.3 | WRED commands (supported only on 7210 SAS-D)..... | 287 |
| 12.7.1.4 | Operational commands..... | 288 |
| 12.7.1.5 | Show commands..... | 288 |
| 12.8 | Command descriptions..... | 288 |
| 12.8.1 | Configuration commands..... | 288 |
| 12.8.1.1 | Generic commands..... | 288 |
| 12.8.2 | Operational commands..... | 289 |
| 12.8.2.1 | Slope policy QoS commands..... | 289 |
| 12.8.2.2 | Slope policy QoS policy commands..... | 290 |
| 12.8.2.3 | RED slope commands..... | 293 |
| 12.8.2.4 | WRED command for 7210 SAS-D..... | 296 |
| 12.8.3 | Show commands..... | 298 |
| | slope-policy..... | 298 |
| 13 | Standards and protocol support..... | 303 |
| 13.1 | BGP..... | 303 |
| 13.2 | Ethernet..... | 305 |
| 13.3 | EVPN..... | 306 |
| 13.4 | Fast Reroute..... | 306 |
| 13.5 | Internet Protocol (IP) — General..... | 307 |
| 13.6 | IP — Multicast..... | 309 |

| | | |
|-------|-------------------------|-----|
| 13.7 | IP — Version 4..... | 310 |
| 13.8 | IP — Version 6..... | 311 |
| 13.9 | IPsec..... | 312 |
| 13.10 | IS-IS..... | 313 |
| 13.11 | Management..... | 314 |
| 13.12 | MPLS — General..... | 317 |
| 13.13 | MPLS — GMPLS..... | 318 |
| 13.14 | MPLS — LDP..... | 318 |
| 13.15 | MPLS — MPLS-TP..... | 318 |
| 13.16 | MPLS — OAM..... | 319 |
| 13.17 | MPLS — RSVP-TE..... | 319 |
| 13.18 | OSPF..... | 320 |
| 13.19 | Pseudowire..... | 321 |
| 13.20 | Quality of Service..... | 322 |
| 13.21 | RIP..... | 322 |
| 13.22 | Timing..... | 322 |
| 13.23 | VPLS..... | 324 |

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-D and 7210 SAS-Dxp.....

26

Table 6: Default network QoS policy egress marking on 7210 SAS-D.....

28

Table 7: Default network QoS policy egress marking on 7210 SAS-Dxp.....

29

Table 8: Default network QoS policy ingress classification.....

30

Table 9: Default network queue policy definition (7210 SAS-D).....

30

Table 10: Default network queue policy definition (7210 SAS-Dxp).....

31

Table 11: Default service ingress policy ID 1 definition.....

33

Table 12: Default access egress policy ID 1 definition for 7210 SAS-D and 7210 SAS-Dxp.....

36

Table 13: Supported hardware rates and burst step sizes for CIR and PIR values for 7210 SAS-Dxp.....

39

Table 14: Forwarding classes.....

43

Table 15: FC-to-queue ID map.....

44

Table 16: Minimum and maximum bandwidth shapers example.....

45

Table 17: Supported hardware rates and CIR and PIR values for egress queues on the 7210 SAS-D.....

49

Table 18: Supported hardware rates and CIR and PIR values for egress queues on the 7210 SAS-Dxp....

49

Table 19: Default CBS and MBS values.....

51

Table 20: TAF impact on shared buffer average utilization calculation.....

57

Table 21: Default slope policy definition for 7210 SAS-D.....

59

| | |
|--|-----|
| Table 22: Default slope policy definition for 7210 SAS-Dxp..... | 60 |
| Table 23: Output fields: show PoE port (Ethernet)..... | 69 |
| Table 24: Output fields: QoS SAP-ingress..... | 82 |
| Table 25: Output fields: show QoS network..... | 84 |
| Table 26: Output fields: access egress..... | 87 |
| Table 27: Output fields: network queue..... | 89 |
| Table 28: Output fields: port scheduler policy..... | 90 |
| Table 29: Default network policy 1..... | 102 |
| Table 30: DSCP and dot1p marking..... | 102 |
| Table 31: Default DSCP mapping table..... | 103 |
| Table 32: Output fields: QoS network..... | 136 |
| Table 33: Network queue policy defaults for 7210 SAS-D..... | 139 |
| Table 34: Network queue policy defaults for 7210 SAS-Dxp..... | 141 |
| Table 35: Output fields: network queue policy..... | 154 |
| Table 36: SAP ingress policy defaults..... | 156 |
| Table 37: SAP ingress resource allocation and match criteria types..... | 161 |
| Table 38: DSCP mask value format..... | 215 |
| Table 39: IP precedence mask value format..... | 218 |
| Table 40: 3-bit mask format..... | 221 |
| Table 41: dot1p mask value format..... | 222 |
| Table 42: MAC mask format..... | 224 |
| Table 43: Output fields: QoS SAP ingress..... | 236 |
| Table 44: Default FC marking values for 7210 SAS-D and 7210 SAS-Dxp..... | 241 |

| | |
|--|-----|
| Table 45: Access-egress labels and descriptions..... | 258 |
| Table 46: Output fields: show port scheduler policy..... | 269 |
| Table 47: Slope behavior (applicable to 7210 SAS-D)..... | 270 |
| Table 48: Slope behavior (applicable to 7210 SAS-Dxp)..... | 271 |
| Table 49: Slope policy defaults for 7210 SAS-D..... | 274 |
| Table 50: Slope policy defaults for 7210 SAS-Dxp..... | 278 |
| Table 51: Output fields: slope policy..... | 299 |
| Table 52: Output fields: slope policy detail..... | 301 |

List of figures

Figure 1: 7210 SAS-D and 7210 SAS-Dxp service and network traffic types and QoS model.....27

Figure 2: Traffic policing and queuing model for FCs.....33

Figure 3: RED slope characteristics.....56

Figure 4: Calculation for average shared buffer utilization..... 57

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

This guide describes the Quality of Service (QoS) functionality provided by the following 7210 SAS platforms, operating in one of the modes listed in [Table 1: Supported modes of operation and configuration methods](#). If multiple modes of operation apply, they are explicitly noted in the topic.



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.

- 7210 SAS-D
- 7210 SAS-Dxp 12p (2SFP+ 4SFP 6Tx)
- 7210 SAS-Dxp 16p (2SFP+ 4SFP 10Tx)
- 7210 SAS-Dxp 24p (2SFP+ 6SFP 16Tx)

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 described in this document" is used to indicate that the topic and CLI apply to all the following 7210 SAS platforms implicitly operating in the specified modes only. See [Table 1: Supported modes of operation and configuration methods](#) for more information.

- **access-uplink mode of operation**
7210 SAS-D and 7210 SAS-Dxp
- **standalone mode of operation**
7210 SAS-D and 7210 SAS-Dxp

1.1.1 Document structure and content

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



Note:

This guide generically covers Release 25.x.Rx content and may include some content that will be released in later maintenance loads. See the *7210 SAS Software Release Notes 25.x.Rx*, part

number 3HE 21188 000x TQZZA, for information about features supported in each load of the Release 25.x.Rx software.

- This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow. Each chapter describes a software area and provides CLI syntax and command usage to configure parameters for the functional area.
- Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.

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 25.x.Rx*, part number 3HE 21188 000x TQZZA, and to the appropriate 7210 SAS software user guide for information about features and capabilities supported by a 7210 SAS platform when operating in a specific mode.

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

- **access-uplink**

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

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

- **network**

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

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

- **satellite**

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

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

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

- **standalone**

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

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

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

- **standalone-VC**

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

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

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

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

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

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

Table 1: Supported modes of operation and configuration methods

| 7210 SAS platform | Mode of operation and configuration method | | | | |
|--------------------------------|--|--------------------------------------|------------|---------------|-----------|
| | Network | Access-uplink | Standalone | Standalone-VC | Satellite |
| 7210 SAS-D | | Implicit | Implicit | | |
| 7210 SAS-Dxp | | Implicit | Implicit | | |
| 7210 SAS-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 ⁴ | | Explicit BOF Configuration | | Explicit BOF Configuration |
| 7210 SAS-R6 ³ | Implicit | | Implicit | | |
| 7210 SAS-R12 ³ | Implicit | | Implicit | | |
| 7210 SAS-Sx/S 1/10GE | Implicit ⁴ | | Explicit BOF Configuration | Explicit BOF Configuration | Explicit BOF Configuration |
| 7210 SAS-Sx 10/100GE | Implicit ⁴ | | 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 Supports MPLS uplinks only and implicitly operates in network mode

⁴ Implicitly operates in [network](#) mode when [standalone](#) mode of operation is configured

- **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.

[Table 2: Supported port modes by mode of operation](#) 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 |

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-K 3SFP+ 8C | Yes | Yes | Yes | Yes |
| 7210 SAS-Mxp | Yes | Yes | Yes | No |
| 7210 SAS-R6 IMM-b (IMMv2) | Yes | Yes | Yes | No |
| 7210 SAS-R6 IMM-c 100GE (IMM-c 1QSFP28) | Yes | Yes | Yes | No |
| 7210 SAS-R12 IMM-b | Yes | Yes | Yes | No |
| 7210 SAS-R12 IMM-c 100GE (IMM-c 1QSFP28) | Yes | Yes | Yes | No |
| 7210 SAS-Sx/S 1/10GE | Yes | Yes | Yes | No |
| 7210 SAS-Sx 10/100GE | Yes | Yes | Yes | No |
| 7210 SAS-T | Yes | Yes ⁶ | Yes ⁷ | Yes ⁸ |

1.4 Nokia 7210 SAS-series services configuration process

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

Table 4: Configuration process

| Area | Task | Chapter |
|----------------------|--------------------------|---|
| Policy configuration | Configuring QoS policies | |
| | • egress rate | Port level egress rate-limiting |
| | • accounting mode | Frame-based accounting |
| | • network | Network QoS policies |
| | • network queue | Network queue QoS 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 |
|-----------|--|--|
| | • SAP ingress | Service ingress QoS policies |
| | • access egress | Access egress QoS policies |
| | • port scheduler | Port scheduler policies |
| | • slope | Slope QoS policies |
| Reference | • list of IEEE, IETF, and other proprietary entities | Standards and protocol support |

1.5 Conventions

This section describes the general conventions used in this guide.

1.5.1 Precautionary and information messages

The following information symbols are used in the documentation.



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



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



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



Note: Note provides additional operational information.



Tip: Tip provides suggestions for use or best practices.

1.5.2 Options or substeps in procedures and sequential workflows

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

Example: Options in a procedure

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

- This is one option.
- This is another option.
- This is yet another option.

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

Example: Substeps in a procedure

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

2 QoS policies

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

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, police, queue, shape, and mark traffic.



Note:

Not all QoS capabilities are supported on all 7210 SAS platforms. 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. 7210 SAS-D and 7210 SAS-Dxp support QinQ uplinks or dot1q uplinks or NULL port for transport of services.

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

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 on 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" are 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 types:

- QoS 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.

2.2 Overview of QoS policies on 7210 SAS-D and 7210 SAS-Dxp

7210 SAS-D and 7210 SAS-Dxp QoS policies are applied on service ingress, access port egress, and access uplink ports (ingress and egress). These policies allow users to configure the following:

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

There are the following types of QoS policies:

- service ingress policies for access SAP ingress
- access egress policies for access port egress
- network policies for access-uplink port ingress and egress
- network queue policies for access-uplink port egress
- port scheduler for access port and access-uplink port egress
- slope policies for congestion management using RED

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

A service ingress QoS policy also defines up to three meters per FC, which can be used for multipoint traffic for multipoint services. There can be up to 32 meters in total per service ingress QoS policy.

For VPLS, the following types of forwarding are supported:

- unicast
- multicast
- broadcast
- unknown

Multicast, broadcast, and unknown forwarding types are typically sent to multiple destinations within the service, while the unicast forwarding type is handled in a point-to-point manner within the service.

An access egress policy is analogous to a SAP egress policy. The difference is the point of attachment. An access-egress policy is applied on the physical port as opposed to the logical port (SAP) for SAP egress policy and applies to the traffic sent out of all the SAPs configured on the port. An access-egress QoS policy maps the traffic egressing out on the customer facing ports into various queues and marks the traffic accordingly. The FCs are mapped onto the queues with an option to configure the queue parameters (for example, rate values). There are eight egress queues at the port level. FC-to-queue mapping is static and is not configurable. The number of queues are static and there are always eight egress queues at the port level. An access-egress policy also defines how to remark the FC-to-priority bits (for example, IEEE 802.1p bits) in the customer traffic.

Network QoS policies can be applied to access uplink ports. On ingress, the policy can be used to map incoming dot1p values to the FC and profile state for the traffic received from the core network. On egress,

the policy maps the FC and profile state to priority bit (for example, IEEE 802.1p bits) values for traffic to be transmitted into the core network.

On egress, network queue policies are applied to access uplink ports. The policies define the FC queue characteristics.

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

One service ingress QoS policy can be applied to a specific SAP access egress policy, which can then be applied to an access port, with a single access egress QoS policy allowed to be associated with an access port. One network QoS policy can be applied to a specific access-uplink port. A network QoS policy defines both ingress and egress behavior. 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.

2.2.1 Summary of major functions of QoS policies

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

Table 5: QoS policy types and descriptions for 7210 SAS-D and 7210 SAS-Dxp

| Policy type | Applied at... | Description |
|-----------------|------------------------------------|---|
| Service ingress | Access SAP ingress | <ul style="list-style-type: none">defines up to 32 FC meters and meter parameters for traffic classificationdefines match criteria to map flows to the meters based on any one of the criteria (IP or MAC or both IP and MAC) |
| Access egress | Access port | <ul style="list-style-type: none">defines up to eight FC queues and queue parameters for traffic classificationmaps FCs to the queuesdefines queue parameters for the queuesdefines FC-to-remarking valuesdefines CIR levels and PIR weights that determine how the queue is prioritized by the scheduler |
| Egress rate | Access port and access-uplink port | <ul style="list-style-type: none">configures the maximum bandwidth available for traffic sent out of a specified port |
| Accounting mode | Device level | <ul style="list-style-type: none">sets the accounting mode to packet-based or frame-based for ingress and egress QoS policies |
| Network | Access uplink ports | <ul style="list-style-type: none">on ingress, defines dot1p to FC mapping and meterson egress, defines FC-to-remarking values (for example, IEEE 802.1p) |


| Policy type | Applied at... | Description |
|----------------|------------------------------------|---|
| Network queue | Access uplink ports | <ul style="list-style-type: none"> defines FC mappings to network queues and queue characteristics for the queues |
| Slope | Port queues | <ul style="list-style-type: none"> on 7210 SAS-D, enables or disables the high-slope, low-slope, and non-TCP parameters within the egress pool on 7210 SAS-Dxp, enables or disables high-slope and low-slope parameters for TCP traffic |
| Port scheduler | Access port and Access-uplink port | <ul style="list-style-type: none"> defines the parameters for the port scheduler |

2.3 Network and service QoS policies

The QoS mechanisms within the 7210 SAS-D and 7210 SAS-Dxp are specialized to cater to the type of traffic on the interface.

The following figure shows that on 7210 SAS-D and 7210 SAS-Dxp, for customer interfaces, there are service ingress and access port egress traffic types, and for access uplink port interfaces, there are network ingress and network egress traffic types.

Figure 1: 7210 SAS-D and 7210 SAS-Dxp service and network traffic types and QoS model



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

The 7210 SAS supports the following major types of service and network QoS policies: [Network QoS policies on access-uplink ports](#), [Default network QoS policy \(egress\) for the 7210 SAS-Dxp](#), [Default network QoS policy \(ingress\)](#), [Network queue policies in access-uplink mode](#), [Service ingress QoS policies](#), and [Service ingress classification](#).

2.3.1 Network QoS policies on access-uplink ports

Network QoS policies define egress QoS marking and ingress QoS interpretation for traffic received on access-uplink ports. The router automatically creates egress queues for each of the FCs on the access-uplink port.

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

- **ingress**
 - defines dot1p value mapping to FCs and profile
 - defines FC to meter mapping
- **egress**
 - defines 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 and 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 or out-of-profile state
- at least one default unicast FC meter. See [Committed information rate \(meters\)](#) for information about meter configuration parameters.
- optional multipoint FC meter

The optional network QoS policy elements definitions are the following:

- additional unicast meters or queues
- additional multipoint meters or 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 define the FC and profile state to IP DSCP value marking

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

The network QoS policy applied at network egress (that is, on an access-uplink port egress) 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 7210 SAS-D and 7210 SAS-Dxp. The following table lists the default mapping of FC to dot1p values for egress marking on the 7210 SAS-D.

Table 6: Default network QoS policy egress marking on 7210 SAS-D

| FC-ID | FC name | FC label | DiffServ name | Egress dot1p marking | | Egress DSCP marking | |
|-------|-----------------|----------|---------------|----------------------|----------------|---------------------|----------------|
| | | | | In-profile | Out-of-profile | In-profile | Out-of-profile |
| 7 | Network Control | nc | NC2 | 111 - 7 | 111 - 7 | nc2 | nc2 |

| FC-ID | FC name | FC label | DiffServ name | Egress dot1p marking | | Egress DSCP marking | |
|-------|-------------|----------|---------------|----------------------|----------------|---------------------|----------------|
| | | | | In-profile | Out-of-profile | In-profile | Out-of-profile |
| 6 | High-1 | h1 | NC1 | 110 - 6 | 110 - 6 | nc1 | nc1 |
| 5 | Expedited | ef | EF | 101 - 5 | 101 - 5 | ef | ef |
| 4 | High-2 | h2 | AF4 | 100 - 4 | 100 - 4 | af41 | af41 |
| 3 | Low-1 | l1 | AF2 | 011 - 3 | 011 - 3 | af21 | af22 |
| 2 | Assured | af | AF1 | 011-3 | 011-3 | af11 | af12 |
| 1 | Low-2 | l2 | CS1 | 001 - 1 | 001 - 1 | cs1 | cs1 |
| 0 | Best Effort | be | BE | 000 - 0 | 000 - 0 | be | be |

2.3.2 Default network QoS policy (egress) for the 7210 SAS-Dxp

The following table lists the default mapping of FC to dot1p values for egress marking on the 7210 SAS-Dxp.

Table 7: Default network QoS policy egress marking on 7210 SAS-Dxp

| FC-ID | FC name | FC label | DiffServ name | Egress dot1p marking | | Egress DSCP marking | |
|-------|-----------------|----------|---------------|----------------------|----------------|---------------------|----------------|
| | | | | In-profile | Out-of-profile | In-profile | Out-of-profile |
| 7 | Network Control | nc | NC2 | 111 - 7 | 111 - 7 | nc2 | nc2 |
| 6 | High-1 | h1 | NC1 | 110 - 6 | 110 - 6 | nc1 | nc1 |
| 5 | Expedited | ef | EF | 101 - 5 | 101 - 5 | ef | ef |
| 4 | High-2 | h2 | AF4 | 100 - 4 | 100 - 4 | af41 | af41 |
| 3 | Low-1 | l1 | AF2 | 011 - 3 | 011 - 3 | af21 | af22 |
| 2 | Assured | af | AF1 | 011-3 | 011-3 | af11 | af12 |
| 1 | Low-2 | l2 | CS1 | 001 - 1 | 001 - 1 | cs1 | cs1 |
| 0 | Best Effort | be | BE | 000 - 0 | 000 - 0 | be | be |

2.3.3 Default network QoS policy (ingress)

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 8: Default network QoS policy ingress classification

| dot1p value | 7210 FC ingress | Profile |
|-------------|-----------------|---------|
| 0 | be | Out |
| 1 | l2 | In |
| 2 | af | Out |
| 3 | af | In |
| 4 | h2 | In |
| 5 | ef | In |
| 6 | h1 | In |
| 7 | nc | In |

2.3.4 Network queue policies in access-uplink mode

In access-uplink mode of operation, network queue policies are applied on egress of access-uplink ports.

On 7210 SAS-D and 7210 SAS-Dxp, the system allocates eight egress queues (not user-configurable) for the network port, and FCs are mapped to these eight egress queues. All policies use eight egress queues like the default network queue policy. The following table describes the eight egress queues for 7210 SAS-D.

Table 9: Default network queue policy definition (7210 SAS-D)

| FC | Queue | Definition |
|----------------------|---------|---|
| Network-Control (nc) | Queue 8 | <ul style="list-style-type: none">PIR = 100%CIR = 10% |
| High-1 (h1) | Queue 7 | <ul style="list-style-type: none">PIR = 100%CIR = 10% |
| Expedited (ef) | Queue 6 | <ul style="list-style-type: none">PIR = 100%CIR = 100% |
| High-2 | Queue 5 | <ul style="list-style-type: none">PIR = 100%CIR = 100% |

| FC | Queue | Definition |
|-------------|---------|---|
| Low-1 | Queue 4 | <ul style="list-style-type: none"> PIR = 100% CIR = 25% |
| Assured | Queue 3 | <ul style="list-style-type: none"> PIR = 100% CIR = 25% |
| Low-2 | Queue 2 | <ul style="list-style-type: none"> PIR = 100% CIR = 25% |
| Best Effort | Queue 1 | <ul style="list-style-type: none"> PIR = 100% CIR = 0% |

The following table describes the eight egress queues for 7210 SAS-Dxp.

Table 10: Default network queue policy definition (7210 SAS-Dxp)

| FC | Queue | Definition |
|----------------------|---------|---|
| Network-Control (nc) | Queue 8 | <ul style="list-style-type: none"> PIR = 100% CIR = 10% |
| High-1 (h1) | Queue 7 | <ul style="list-style-type: none"> PIR = 100% CIR = 5% |
| Expedited (ef) | Queue 6 | <ul style="list-style-type: none"> PIR = 100% CIR = 15% |
| High-2 | Queue 5 | <ul style="list-style-type: none"> PIR = 100% CIR = 15% |
| Low-1 | Queue 4 | <ul style="list-style-type: none"> PIR = 100% CIR = 10% |
| Assured | Queue 3 | <ul style="list-style-type: none"> PIR = 100% CIR = 10% |
| Low-2 | Queue 2 | <ul style="list-style-type: none"> PIR = 100% CIR = 10% |
| Best Effort | Queue 1 | <ul style="list-style-type: none"> PIR = 100% CIR = 0% |

2.3.5 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.



Note:

Not all 7210 platforms support queues and meters on service ingress. The support varies across different platforms. Please read the subsequent chapters and sections for more information.

When a service ingress QoS policy is created, it typically has some meters defined that cannot be deleted and is used for all traffic (both unicast and multicast traffic). These meters exist within the definition of the policy but only get instantiated in hardware when the policy is applied to a SAP. In a case where the service does not have multipoint traffic, for example Epipe service, the multipoint meters are not instantiated.

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

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

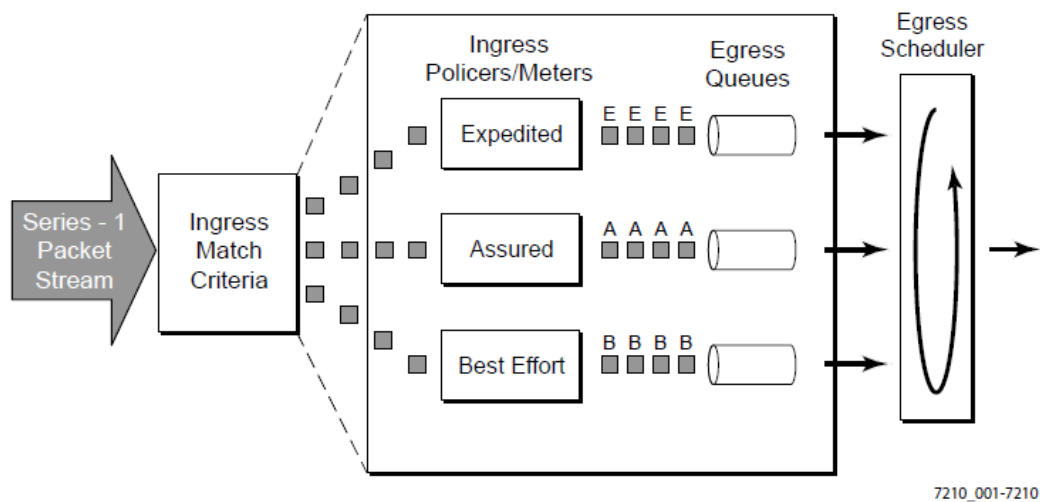
- a unique service ingress QoS policy ID
- a QoS policy scope of **template** or **exclusive**
- the number of classification and meter resources to allocate for this policy
- the allocation of resources from the ingress internal CAM resource pool for use with service ingress QoS classification using the commands available under the CLI context **configure>system>resource-profile**. Additionally, the allocation of resources to the appropriate classification match criteria
- at least one default FC meter. See [Metering/policing and meter parameters](#) for information about the parameters that can be configured for a meter.

Optional service ingress QoS policy elements include the following:

- additional unicast meters, up to 31
- additional multipoint meters up to 31
- QoS policy match criteria to map packets to an FC

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

Figure 2: Traffic policing and queuing model for FCs



2.3.5.1 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 have another service ingress policy assigned. In the default policy, all traffic is mapped to the default FC, which uses a meter by default. The following table lists the characteristics of the default policy.

Table 11: Default service ingress policy ID 1 definition

| Item | Definition |
|-----------------|---|
| Meter 1 | 1 (one) meter all unicast traffic: <ul style="list-style-type: none">FC: best-effort (be)CIR = 0PIR = max (4000000 kb/s in case of a LAG with four member ports)CBS = 32 kbMBS = 128 kb |
| Default FC (be) | 1 (one) flow defined for all traffic: <ul style="list-style-type: none">all traffic mapped to best-effort (be) |

2.3.6 Service ingress classification

Mapping flows to FCs are 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 provisioned classification policy.

2.3.6.1 Service ingress classification

On SAP ingress, users have an option to use either MAC criteria or IP criteria, or both IPv4 and MAC criteria. To use the available classification resources effectively, the following options are available:

- supported MAC header fields using the **mac-criteria any** option
- only dot1p bits in the MAC header using the **mac-criteria dot1p-only** option
- supported IPv4 header fields using the **ip-criteria any** option
- only IPv4 DSCP in the IPv4 header using the **ip-criteria dscp-only** option
- supported IPv6 header fields using the **ipv6-criteria any** option
- only IPv6 DSCP bits in the IPv6 header using the **ipv6-criteria dscp-only** option
- both MAC and IPv4 header fields using the both MAC and IPv4 criteria option together in a policy (this option is supported only on 7210 SAS-D and 7210 SAS-Dxp)

Among the above supported criteria the following can be configured together in a single policy:

- **mac-criteria any**
- **mac-criteria dot1p-only**
- **ip-criteria any** and **ipv6-criteria any** or **ipv6-criteria dscp-only**
- **ipv6-criteria dscp-only** and **ipv6-criteria any** or **ipv6-criteria dscp-only**
- **mac-criteria any** and **ip-criteria any** or **ip-criteria dscp-only** or **ipv6-criteria dscp-only** (this option is supported only on 7210 SAS-D and 7210 SAS-Dxp)
- **mac-criteria dot1p-only** and **ip-criteria any** or **ip-criteria dscp-only** or **ipv6-criteria dscp-only** (this option is supported only on 7210 SAS-D and 7210 SAS-Dxp)



Note:

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

The available ingress CAM hardware resources from the ingress internal CAM pool can be allocated according to user needs for use with different QoS classification match criteria using the commands available under the CLI **configure>system> resource-profile** context. By default, the system allocates resources to SAP ingress classification on bootup. Users can change the resource allocation based on their need to scale the number of entries or number of associations (that is, the number of SAPs using a policy that uses a particular match criterion). If no resources are allocated to a particular match criteria used in the policy, the association of that policy to a SAP fails. Allocation of classification entries also allocates meter resources, used to implement the per FC per traffic type policing function. See [Resource allocation for service ingress QoS policy classification rules](#) for more information about resource usage and allocation to SAP ingress policies.

In addition to the preceding classification rules, the user has an option to use the DEI bit to identify the ingress profile and enable color-aware policing. See [DEI-based classification and marking](#) for more information.

2.3.6.2 Hierarchical ingress policing on 7210 SAS-D and 7210 SAS-Dxp

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

Hierarchical ingress policing provides an option to configure a SAP aggregate policer/meter per SAP on SAP ingress. The user should configure the PIR rate of the aggregate policer and optionally configure the burst size.

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

- per FC policer
- per SAP aggregate policer

For more information about the final color/profile assigned to the packet, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Services Guide*, **configure service sap ingress aggregate-meter-rate** command description.

The **trtcm2** (RFC 4115) meter mode is used only on SAP ingress. When the SAP aggregate policer is configured, the per-FC policer can be only configured in "trtcm2" mode. The meter modes "srtcm" and "trtcm1" can be used in the absence of an aggregate meter.



Note:

Before using per-SAP aggregate meter/policer, meter resources must be allocated using the **config system resource-profile ingress-internal-tcam sap-aggregate-meter** command. These resources are shared with ingress ACLs. A change to the number of resources allocated for the SAP aggregate meter requires a reboot of the node to take effect. The amount of resources allocated for this feature determines the number of SAPs that can use the aggregate meter/policer. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for more information.

2.4 Access egress QoS policies

On the 7210 SAS-D and 7210 SAS-Dxp, an access-egress policy defines the queue and marking characteristics for the traffic egressing toward the customer on the access ports. There are eight queues always available at the access port, and FCs are mapped to these eight queues. By configuring queue shaper rates, the individual FC traffic can be managed so that the traffic for each FC is well within SLA limits and does not impact the traffic of other FCs. The access-egress policy also provides an option for marking packets sent out of access ports, allowing the FC to be mapped to packet header priority bits (for example, IEEE dot1p bits).

The FCs are mapped to eight queues by the software, as per [Table 15: FC-to-queue ID map](#), and this mapping is not user configurable. The queue ID determines the priority of the queue, with a higher queue-ID denoting a higher priority.

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

- a unique service access QoS policy ID
- a QoS policy scope of **template** or **exclusive**
- IEEE 802.1p priority value remarking based on FC
- on 7210 SAS-D and 7210 SAS-Dxp, an option is available to use IP DSCP values for marking based on FC

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

The FC determination per service egress packet is determined at ingress. If the packet ingressed the service on the same router, the service ingress classification rules determine the FC of the packet. If the packet was received over a service transport tunnel on an access-uplink port, the FC is marked in the outer tag of the QinQ encapsulation.

2.4.1 Default access egress policy

Access-egress QoS policy ID 1 is reserved as the default access-egress policy associated with access ports that do not have another access-egress policy explicitly assigned. The characteristics of the default policy are listed in [Table 12: Default access egress policy ID 1 definition for 7210 SAS-D and 7210 SAS-Dxp](#).

Table 12: Default access egress policy ID 1 definition for 7210 SAS-D and 7210 SAS-Dxp

| FC | Queue-ID | Queue parameters |
|----------------------|-----------|---|
| Queues | Queue 1-8 | 1 (one) queue defined for each traffic class |
| Network-Control (nc) | Queue 8 | <ul style="list-style-type: none">• CIR=0• PIR= max (line rate) |
| High-1 (h1) | Queue7 | <ul style="list-style-type: none">• CIR=0• PIR= max (line rate) |
| Expedited (ef) | Queue 6 | <ul style="list-style-type: none">• CIR = 0• PIR = max (line rate) |
| High-2 (h2) | Queue 5 | <ul style="list-style-type: none">• CIR = 0• PIR = max (line rate) |
| Low-1 (l1) | Queue 4 | <ul style="list-style-type: none">• CIR = 0• PIR = max (line rate) |
| Assured (af) | Queue 3 | <ul style="list-style-type: none">• CIR = 0• PIR = max (line rate) |
| Low-2 (l2) | Queue 2 | <ul style="list-style-type: none">• CIR = 0 |

| FC | Queue-ID | Queue parameters |
|------------------|----------|--|
| | | <ul style="list-style-type: none"> PIR = max (line rate) |
| Best-Effort (be) | Queue 1 | <ul style="list-style-type: none"> CIR = 0 PIR = max (line rate) |

2.5 Meters/policers

This section provides information about meters/policers.

2.5.1 Metering/policing and meter parameters

This section describes the meter behavior and available meter parameters that can be defined for meters provisioned on the service entities (for example, access SAP ingress on 7210 SAS-D).



Note:

Not all 7210 SAS platforms support meters for all the policies. In addition, the meter parameters supported vary across platforms. See the platform-specific QoS overview sections in this guide. In the following sections, the platform differences about support are noted.

2.5.1.1 Meter ID

The meter ID is used to uniquely identify the meter. The meter ID is only unique within the context of the QoS policy in which the meter is defined. It allows users to define multiple meters in a policy and associate them with one of the eight FCs.

2.5.1.2 Committed information rate (meters)

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 can configure the burst above the CIR and up to the PIR for brief periods of time. The amount of burst is determined by the CBS and MBS values configured for the meter.

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 some control over how the administrative CIR is converted to an operational CIR if the hardware does not support the exact CIR and PIR combination specified. See [Adaptation rule for meters](#) for more information about the interpretation of the administrative CIR.

2.5.1.3 Peak information rate (meters)

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 hardware that determine the operational PIR for the meter. The user has some control over how the administrative PIR is converted to an operational PIR if the hardware does not support the exact CIR and PIR combination specified. See [Adaptation rule for meters](#) for more information about the interpretation of the administrative PIR.

2.5.1.4 Adaptation rule for meters

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

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

2.5.1.4.1 Adaptation rule for meters on 7210 SAS-D devices

The hardware supports meter rates in the multiples of 8 kb/s for the entire range of CIR or PIR rates supported on the device. The system identifies the best operational rate depending on the defined constraint. The following constraints are supported:

- **minimum**
The system finds the next multiple of 8 kb/s that is equal to or greater than the specified rate.
- **maximum**
The system finds the next multiple of 8 kb/s that is less than or equal to the specified rate.
- **closest**
The system finds the next multiple of 8 kb/s that is closest to the specified rate.

2.5.1.4.2 Adaptation rule for meters on 7210 SAS-Dxp devices

The hardware supports meter rates in the multiples of 8 kb/s for CIR or PIR rates on 1G ports supported on the device. The system identifies the best operational rate depending on the defined constraint. The following constraints are supported for 1G ports:

- **minimum**
The system finds the next multiple of 8 kb/s that is equal to or greater than the specified rate.
- **maximum**

The system finds the next multiple of 8 kb/s that is less than or equal to the specified rate.

- **closest**

The system finds the next multiple of 8 kb/s that is closest to the specified rate.

The following table lists information for calculating the hardware-supported meter step size for all supported ranges of rate values and burst step sizes for all supported ranges of burst values for 10G ports.

Table 13: Supported hardware rates and burst step sizes for CIR and PIR values for 7210 SAS-Dxp

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



Note:

The configured burst size affects the rate step-size used by the system. The system uses the step size to ensure that both the burst size and rate parameter constraints are met. For example, if the rate specified is less than 4 Gb/s but the configured burst size is 17 Mb, then the system uses rate step size of 16 Kb and burst step size of 8192 bits (see [Table 13: Supported hardware rates and burst step sizes for CIR and PIR values for 7210 SAS-Dxp](#), row 2).

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

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

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

2.5.1.5 Committed burst size (meters/policers)

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.



Note:

See [Table 13: Supported hardware rates and burst step sizes for CIR and PIR values for 7210 SAS-Dxp](#) for information about the burst parameter step size for 7210 SAS-Dxp.

2.5.1.6 Maximum burst size (meters/policers)

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.



Note:

See [Table 13: Supported hardware rates and burst step sizes for CIR and PIR values for 7210 SAS-Dxp](#) for information about the burst parameter step size for 7210 SAS-Dxp.

2.5.1.6.1 Excess burst size (meters/policers)

The excess burst size (EBS) parameter specifies the excess burst size transmitted by the source while not complying with the CIR. If the transmitted burst is lower than the EBS value, the packets are marked as out-profile by the meter to indicate that the traffic is not complying with the CIR. If the packet burst is higher than EBS, packets marked as red are dropped.



Note:

- The EBS parameter is used only when the meter mode is set to "trtcm2."
- See [Table 13: Supported hardware rates and burst step sizes for CIR and PIR values for 7210 SAS-Dxp](#) for information about the burst parameter step size for 7210 SAS-Dxp.

2.5.2 Meter counters

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

- counters for packets and octets marked as in-profile by meter
- counters for packets and octets marked as out-of-profile by meter
- counters for packets and octets marked as dropped by meter

The counters available vary among the 7210 SAS platforms. Not all counters are supported on all platforms. Additionally, there are restrictions on the number of counters that can be used simultaneously with a single meter. Some platforms can only count octets or packets and others can count both packets and octets. Typically, each meter can maintain a subset of the counters. Users have the option to select the subset of the counter they want to use. See "Accounting Records" in the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C System Management Guide for information about counter (and corresponding accounting record) support available on each of the platforms.

2.5.3 Meter modes

The meter mode command allows users to select from the following meter types:

- **srTCM**

Single Rate Three Color meter, with this meter a single rate can be specified by the user to limit the amount of traffic. The single rate along with the CBS and MBS determine the amount of both in-profile/committed traffic and out-of-profile/excess burst. With this meter, the CBS and MBS token buckets are replenished at single rate, that is, CIR.

- **trTCM**

Two Rate Three Color meter, with this meter users can specify two rates; CIR and PIR can be specified by the user to limit the amount of traffic. It allows the user to limit the amount of in-profile/committed traffic to CIR rate and allows user to send excess out-of-profile traffic up to a peak rate of PIR. With this meter, the meters CBS token buckets are replenished at the CIR rate and the MBS/EBS token buckets are replenished at the PIR/EIR rate. There are two modes supported under this:

- trTCM1 - implements policing as per RFC 2698
- trTCM2 - implements policing as per RFC 4115
- srTCM - implements policing as per RFC 2697

The 7210 SAS-D and 7210 SAS-Dxp devices support the following meter modes:

- srTCM - Single Rate Three Color Marking (as per RFC 2697)
- trTCM1 - Two Rate Three Color Marking (as per RFC 2698)
- trTCM2 - Two Rate Three Color Marking (as per RFC 4115)

Different QoS policies support different meter modes; not all the meter modes are supported for all the QoS policies.

2.5.3.1 Color-aware and color-blind policing/metering

In color-aware policing, the user can define the color/profile state (color and profile state are used interchangeably in this section) of the packet using the ingress classification rules. The color (also called the profile state) assigned to the packet at ingress is used by the meter along with the rate configured to determine the final profile state of the packet.

The following applies to color-aware policing:

- If the packet is pre-colored as in-profile (also called green packets), depending on the burst size of the packet, the meter can either mark it in-profile or out-profile. If the packet is within the CBS limit, it is assigned a profile value of in-profile (or color value of green). If the packet exceeds the CBS limit, it is reassigned a profile value of out-profile (or color value of yellow).
- If the packet is pre-colored as out-profile (also called yellow packets), even if the packet burst is less than the current available CBS, it would not be marked as in-profile. Instead it is checked against the MBS/EBS limit directly and is assigned a profile value of out-profile (or color value of yellow) if it is within the MBS/EBS limit.
- If the packet burst is higher than the MBS/EBS, it is marked red and dropped by the meter at ingress.

In color-blind policing, the profile and color assigned to the packet on ingress are ignored and all packets are treated as out-of-profile packets. The CIR and PIR rates configured for the meter determine the final

profile and color for the packet. If the packet is within the CIR, the final profile and color assigned to the packet is in-profile/green, and if the packet exceeds the CIR and is within the PIR, the final profile and color assigned to the packet is out-of-profile/yellow. Packets that exceed the PIR rate are dropped.

The final profile state/color marked by the meter on ingress is used to determine the eligibility of the packet to be enqueued into a buffer at egress (when a slope policy is configured at egress).

The 7210 SAS device supports color-aware policing at network ingress; however, at service ingress a policing option is provided to use either color-aware policing or color-blind policing.

The following support is available on the 7210 SAS-D and 7210 SAS-Dxp:

- With access-uplink ports, ingress classification, the profile state/color can be assigned based on either dot1p bits and or DEI.
- Color-aware policing at access-uplink ports ingress is supported by default; the option to use color-blind meter is not available on access-uplink port ingress.
- With access SAP ingress, a policing option is provided to use either color-aware or color-blind policing.

2.5.4 QoS overrides for meters/policers

The QoS override feature allows the user to override the meter parameters, such as CBS, MBS, rate (CIR and PIR), mode, and adaptation rule (CIR and PIR) in the SAP context. It is only supported for access SAPs. The values are taken from the SAP-ingress policy, when the meter parameter values are not overridden.

Meter override commands are supported on the 7210 SAS-D and 7210 SAS-Dxp platforms.

2.5.4.1 Configuration guidelines for QoS override

The configuration guidelines of QoS override are the following:

- QoS override commands can be used only for the meters or policers defined in the SAP ingress policy.
- QoS override commands are not allowed when the attached policy is of an **exclusive** type.
- QoS override commands are not allowed on mirror destination SAPs.
- QoS override commands are not allowed when ToD policy is attached to the SAP.
- QoS override commands are not supported for ingress and egress QoS policies used with access-uplink SAPs and ports.

2.5.4.2 Configuring meter override parameters

Example

The following is a sample meter override parameter configuration output.

```
*7210SAS>config>service>epipe>sap>ingress# info
-----
      qos 13
      meter-override
        meter 1 create
          mode trtcm2
          adaptation-rule pir max cir max
```

```
        cbs 300
        mbs 200
        rate cir 300 pir 400
    exit
exit
-----
*A:7210SAS>config>service>epipe>sap>ingress#
```

2.6 Forwarding classes

The 7210 SAS supports multiple FCs and class-based queuing, which means that the concept of FCs is common to all the QoS policies.

Each FC (also called Class of Service (CoS)) is important only in relation to the other FCs. An FC provides network elements with 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. The following table describes the FCs supported by the 7210 SAS.

Table 14: Forwarding classes

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



Note:

[Table 14: Forwarding classes](#) lists the default definitions for the FCs. The FC behavior, in terms of ingress marking interpretation and egress marking, can be changed.

2.6.1 FC-to-queue ID map

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

On these platforms there are only eight queues available at the port level. These eight queues are created by default per port. Users cannot create or delete the queues or the queue ID. Only the queue parameters can be changed. The queue ID is not a configurable entity and queue ID 1 to 8 are, by default, used to identify these eight queues available on the port. The eight queues are available both on the access and access-uplink ports. Queue parameters for these eight queues can be configured as part of the access-egress QoS policy which is applied on the access ports and network-queue policy which is applied on the access-uplink ports.

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

Table 15: FC-to-queue ID map

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

2.7 Schedulers

This section provides information about schedulers.

2.7.1 Port scheduler policies

Port scheduler policies control the traffic behavior on a per-port basis. Associated with each egress port is a set of eight class of service (CoS) queues and a default port scheduler policy named "default". This default policy makes the port behave in strict mode. The default policy cannot be modified. The user can attach another policy to the port to change its scheduling behavior. The scheduler provides the arbitration across the eight CoS queues and is configured in a variety of modes. A major aspect of the arbitration mechanism is the ability to provide minimum and maximum bandwidth guarantees. After the packets are mapped into a CoS queue, they are forwarded/conditioned using one of the schedulers (such as Strict

Priority (SP), Round-Robin (RR), Weighted Round-Robin (WRR), Weighted Deficit Round-Robin (WDRR), (WRR+SP, WDRR+SP)). The traffic shaping aspect is tightly integrated with the scheduler.

2.7.2 Scheduler modes

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

- **Strict priority scheduling across CoS queues**

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

The following table shows that CoS queues 7 and 6 each have a minimum bandwidth specification of 10 Mbps, whereas the remaining QoS queues have a minimum bandwidth specification of 50 Mbps. All CoS queues have a maximum bandwidth specification of 1 Gbps. The goal of these settings is to guarantee the minimum bandwidth settings for each of the queues while also allowing each CoS queue to fully use the egress port capability by having the maximum bandwidth setting at 1 Gbps. The strict priority scheduler mode provides low latency service for CoS queues 6 and 7 while their minimum bandwidth guarantees are being satisfied.

Table 16: Minimum and maximum bandwidth shapers example

| Queue ID | Minimum bandwidth | Maximum bandwidth |
|----------|-------------------|-------------------|
| 7 | 10 Mbps | 1 Gbps |
| 6 | 10 Mbps | 1 Gbps |
| 5 | 50 Mbps | 1 Gbps |
| 4 | 50 Mbps | 1 Gbps |
| 3 | 50 Mbps | 1 Gbps |
| 2 | 50 Mbps | 1 Gbps |
| 1 | 50 Mbps | 1 Gbps |
| 0 | 50 Mbps | 1 Gbps |

- **Round robin scheduling across CoS queues**

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

- **Weighted round robin (WRR)**

In WRR mode, the scheduler provides access to each CoS queue in round robin order. When the scheduler is providing access to a particular queue, it services a configurable number of back-to-back packets before moving on to the subsequent CoS queue. A value of strict is used to designate that a particular queue be considered to be a part of a hybrid Strict + WRR configuration. The values 1 to 15 are used to indicate the number of back-to-back packets to be serviced when the scheduler is servicing a particular CoS queue. If the weight specified is N, but if the number of packets in the queue is lesser than N, the scheduler continues working and moves on to the next backlogged queue. In this mode, with no strict queues configured, the scheduler services the queues in round robin in the CIR loop. The configured weights are not considered in the CIR loop. The weights are used only in the PIR loop.

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

An inherent limitation of the WRR mode is that bandwidth is allocated in terms of packets. WRR works well if the average packet size for each CoS queue flow is known. WDRR aims at addressing this issue.

WDRR provides a bandwidth allocation scheduler mode that takes into account the variably-sized packet issue by maintaining sufficient state information when arbitrating across the CoS queues. In this mode, with no strict queues configured, the scheduler services the queues in round-robin in the CIR loop. The configured weights are not considered in the CIR loop. The weights are used only in the PIR loop. A weight value of 1 to 15 can be configured for each queue. Based on the weights specified, the respective amount of bytes is removed from the queue. A value of 0 is used to designate that a particular queue be considered to be a part of a hybrid Strict + WDRR configuration. If a weight value of 1 is specific for queue 1 and 5 is specific for queue 2, then we will see traffic out of the port in the ratio of 1:5 between the queues (1 and 2), provided no traffic is flowing in the other queues. A weight value of 1 will actually pump out 2 Kbytes from that queue, a value of 5 will pump out 10 Kbytes. Twice of the weight value specific is pumped out.

- **Strict + WRR/WDRR**

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

2.8 CPU queues

This section describes CPU queues and related features.

2.8.1 Overview

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 used for MAC learning (a copy of which is sent to CPU for MAC learning), EFM, CFM/Y.1731, STP, LACP, ICMP, TWAMP, and so on. A set of queues are used to queue packets to the CPU. The number of queues varies per 7210 SAS platform:

- The 7210 SAS-D has 21 queues.

- The 7210 SAS-Dxp has 8 queues.

The packets destined for the CPU are classified internally and are put into the correct queue. These packets are rate-limited to prevent DoS attacks. The software programs the classification entries to identify these packets and assigns appropriate bandwidth and priority to them. It is not configurable by the user.

2.9 Egress port rate limiting

The 7210 SAS supports port egress 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.

2.10 Queue management

This section provides information about QoS queue management.

2.10.1 Queues and queue parameters

This section describes the parameters available for queues used in an access-egress policy and network-queue policy.



Note:

Not all 7210 SAS platforms support queues for all policies. In addition, the queue parameter support available varies across different platforms. See platform-specific QoS overview sections above and the following chapters for information about the support available on different platforms.

2.10.1.1 Queue ID

The queue ID is used to uniquely identify the queue. The queue ID is only unique within the context of the QoS policy where the queue is defined. It allows the user to define multiple queues with different characteristics and identify them while associating it with different FCs.

The software creates eight queues by default with queue ID 1 to 8. The queue ID is automatically assigned to the eight egress queues by the software; it is not configurable. Only some of the queue parameters that determine the queue characteristics are configurable.

2.10.1.2 Committed information rate

The committed information rate (CIR) for a queue performs the following functions:

- **minimum bandwidth guarantees**

The queue CIR setting provides the bandwidth for the 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-profile state of the queue is linked to scheduler prioritizing behavior, as described in the following bullet point.

- **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. See [Port scheduler policies](#) for information about the scheduler behavior on different 7210 SAS platforms.

The in-profile and out-profile state of the queue does not change the packets profile state based on the queue CIR or PIR values. The in-profile and out-profile state of the queue interacts with the scheduler mechanism and provides the minimum and maximum bandwidth guarantees.

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

Although the 7210 SAS is flexible in how the CIR can be configured, there are conventional ranges for the CIR based on the FC of a queue. 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 — example under access port policies, network queue policies, and so on.



Note:

See [Schedulers](#) for information about queue scheduling support on different 7210 SAS platforms.

2.10.1.3 Peak information rate

The peak information rate (PIR) defines the maximum rate at which packets are allowed to exit the queue. The PIR does not specify the maximum rate at which packets can enter the queue; this is governed by the ability of the queue to absorb bursts. The actual transmission rate of an egress queue depends on more than just its PIR. Each queue is competing for transmission bandwidth with other queues. Each queue, PIR, CIR, and the relative priority and weight of the scheduler serving the queue, all combine to affect the ability of a queue 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 when the hardware does not support the exact CIR and PIR values specified. See [Adaptation rule for queues](#) for more information about the interpretation of the administrative PIR.

The PIR of the queue is configurable under the different QoS policies that provide the option to configure the queue parameters — example under access port policies, network queue policies, and so on.



Note:

See [Schedulers](#) for information about queue scheduling support on different 7210 SAS platforms.

2.10.1.4 Adaptation rule for queues

The adaptation rule provides the QoS provisioning system with the ability to adapt defined CIR and PIR administrative rates to the underlying capabilities of the hardware where 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 following constraints are supported:

- **minimum**

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

- **maximum**

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

- **closest**

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

Depending on the platform where the queue is provisioned, the actual operational CIR and PIR settings used by the queue depend on the method the hardware uses to implement and represent the mechanisms that enforce the CIR and PIR rates. The adaptation rule controls the method the system uses to choose the rate step based on the administrative rates defined by the **rate** command.

On 7210 SAS-D devices, for the supported CIR and PIR range values 0 to 1 Gb, the hardware rates are listed in the following table.

Table 17: Supported hardware rates and CIR and PIR values for egress queues on the 7210 SAS-D

| Hardware rate steps (kb/s) | Rate range (kb/s) |
|----------------------------|------------------------|
| | 0 to 16770 kb/s |
| 16 kb/s | 16780 to 33540 kb/s |
| 32 kb/s | 33550 to 67090 kb/s |
| 64 kb/s | 67100 to 134180 kb/s |
| 128 kb/s | 134190 to 268360 kb/s |
| 256 kb/s | 268370 to 536730 kb/s |
| 512 kb/s | 536740 to 1000000 kb/s |

On 7210 SAS-Dxp devices, for supported CIR and PIR range values 0 to 10 Gb, the hardware rates are listed in the following table.

Table 18: Supported hardware rates and CIR and PIR values for egress queues on the 7210 SAS-Dxp

| Hardware rate steps (kb/s) | Rate range (kb/s) |
|----------------------------|-------------------|
| 64 kb/s | 0 to 134144 kb/s |

| Hardware rate steps (kb/s) | Rate range (kb/s) |
|----------------------------|-----------------------|
| 256 kb/s | 134145 kb/s and above |

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

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

2.10.1.5 Committed burst size (queue)

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 have been 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 options to configure the queue parameters; for example, under service ingress and service egress queue policies, access port policies, network queue policies, and so on. The CBS for a queue is specified in kbytes.

For 7210 SAS-D and 7210 SAS-Dxp, the CBS for the queues is not configurable. The CBS value for the queues is set to appropriate default values that take care of specific FC needs in terms of maintaining the differential treatment. The default values used on different ports are listed in [Table 19: Default CBS and MBS values](#).

2.10.1.6 Maximum burst size (queue)

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 will 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 options to configure the queue parameters; for example, under service ingress and service egress queue policies, access port policies, network queue policies, and so on. The MBS for a queue is specified in Kbytes.

On 7210 SAS-D and 7210 SAS-Dxp, the MBS for the queues is not configurable. The MBS value for the queues is set to appropriate default values that take care of specific FC needs in terms of maintaining the differential treatment. The default values used on different ports are listed in [Table 19: Default CBS and MBS values](#).

2.10.1.6.1 Default CBS and MBS for queues

Table 19: Default CBS and MBS values

| Platforms | CBS (Kbytes) | | MBS (Kbytes) | |
|---------------------------------------|---|---|---------------------------------------|-------------------|
| | Network queue and access-uplink queue | Access queue | Network queue and access uplink queue | Access queue |
| 7210 SAS-D | 8.4 | 8.4 | See ⁹ | See ⁹ |
| 7210 SAS-Dxp 12p | 9 | 9 | See ¹⁰ | See ¹⁰ |
| 7210 SAS-Dxp 16p and 7210 SAS-Dxp 24p | 4.5 (for 1 GE ports) 9 (for 10 GE ports) | 4.5 (for 1 GE ports) 9 (for 10 GE ports) | See ¹¹ | See ¹¹ |

⁹ On 7210 SAS-D, the per-port MBS pool mode is used. With it, the maximum MBS per queue, assuming no other queues on the same port have any traffic, is 78 Kbytes. The **show>pool port ID>access-egress** and **show>pool port ID>access-uplink-egress** commands are available to display the values in use depending on the port mode (either access or access-uplink).

¹⁰ On 7210 SAS-Dxp 12p, the per-port MBS pool mode is used. With it, the maximum MBS per queue, assuming no other queues on the same port have any traffic, is 65 Kbytes. The **show>pool port ID>access-egress** and **show>pool port ID>access-uplink-egress** commands are available to display the values in use depending on the port mode (either access or access-uplink).
The 7210 SAS-Dxp 12p also supports a decommissioning feature with per-port MBS pool mode. With it, the per-port MBS pool can be increased by taking away packet buffers from other ports. In this case, the maximum MBS per queue, assuming no other queue has any traffic on that port, depends on the user configuration. For example, assuming one port is decommissioned and its buffers are allocated to port 1/1/1, the maximum MBS per queue on port 1/1/1, assuming no other queues have any traffic, is equivalent to 202 Kbytes.

¹¹ On 7210 SAS-Dxp 16p and 7210 SAS-Dxp 24p, the per-port MBS limit mode is used. With it, the maximum MBS per queue, assuming no other queues on the same port have any traffic, is 150 Kbytes. The maximum MBS limit per port is 200 Kbytes (the CBS buffers available on that port are in addition to the per-port MBS limit). The **show>pool port ID>access-egress** and **show>pool port ID>access-uplink-egress** commands are available to display the values in use depending on the port mode (either access or access-uplink).
The 7210 SAS-Dxp 16p and 7210 SAS-Dxp 24p also support a decommissioning feature with per-port MBS limit mode. With it, the per-port MBS pool can be increased by taking away packet buffers from other ports. In this case, the maximum MBS per queue, assuming no other queue has any traffic on that port, depends on the user configuration. For example, assuming one port is decommissioned and its buffers are allocated to port 1/1/1, the maximum MBS limit per queue on port 1/1/1, assuming no other queues have any traffic, is equivalent to 186 Kbytes. The maximum MBS limit per port is equivalent to 236 Kbytes (the CBS buffers available on that port are in addition to the per-port MBS limit).

2.10.1.7 Queue counters

The 7210 SAS devices maintain counters for queues within the system for the purpose of granular billing and accounting.

Each queue maintains the following counters:

- counters for packets and octets accepted/forwarded into the queue
- counters for packets and octets rejected at the queue

The counters available vary among the 7210 SAS platforms. 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 others can count both packets and octets. See "Accounting Records" in the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C System Management Guide for information about counter (and corresponding accounting record) support available on each of the platforms.

2.10.2 Buffer pools

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

2.10.2.1 Buffer pools

Buffer pools cannot be created or deleted. The egress buffer pools are distributed as network egress buffer pools for access-uplink ports and access egress buffer pools for access ports. Based on the maximum number of ports to be supported for access and network, the total buffer is distributed into the access egress buffer pool and the network egress buffer pool.

The distribution of the buffers into access and network egress pools take care of the buffer requirements at the port level for various queue shaping and scheduling mechanisms and for various packet sizes varying from 64 bytes to jumbo frames. Each port on the system gets an equal portion of the available buffers. From the buffers allocated to a port, each queue gets its CBS amount of buffers. The remaining buffers are allocated toward the shared MBS pool per port. All the queues of the port can use the buffers from the shared MBS pool and it allows the queue to absorb larger bursts if other queues are not bursting simultaneously.

In addition, for 7210 SAS-Dxp in per-port MBS pool mode, an option is provided to decommission the port and allocate its buffers toward other ports.

2.10.2.2 Decommissioning ports with per-port MBS pool on 7210 SAS-Dxp

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

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

services. This mechanism allows operators who use limited number of ports to deploy services to increase the amount of queue buffers allocated to them by decommissioning ports that are not used to deploy any services.

2.10.2.2.1 Using decommission command for buffer allocation on 7210 SAS-Dxp



Note:

Using the **decommission** command for buffer allocation is only supported on the 7210 SAS-Dxp (all variants). For each port receiving reallocated resources from port decommissioning, a maximum of two ports can be decommissioned.

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

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



Note:

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

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

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

2.10.2.2.2 Configuration guidelines for use of decommission commands on 7210 SAS-Dxp

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

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

The user can modify the list of ports or update to the list of ports specified with the **decommission** command (and also **entry** command) when the node is up, but the changes are effected by software only after a reboot of the node.

The software maintains two lists of entries, one is the current list of ports and another which has been modified by the user and takes effect only after the next reboot. These lists can be displayed using the

show command. The configuration file always stores the list of entries as configured by the user, so that, when rebooted, the modified entries and new entries (if any) take effect.

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

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

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

The following configuration sample shows the ports to be decommissioned.

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

**Note:**

The number of ports that a port can borrow buffers from is limited and varies depending on the platform. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for more information about the **decommission** commands.

2.10.3 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.

2.10.4 RED slopes

This section describes the operation and configuration of random early detection (RED) slopes.

2.10.4.1 Operation and configuration of RED slopes for 7210 SAS-D

On the 7210 SAS-D, each queue provides the user with the following options:

- option to configure three slopes per queue - high-priority RED slope, low-priority RED slope, and a non-TCP RED slope
- option to use two slopes per queue - high-priority RED slope and a low-priority RED slope

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

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

When 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, two 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 the low-priority RED slope. When the queue is configured with option to use non-TCP Slope, non-TCP packets are handled by this slope.

2.10.4.2 Operation and configuration of RED slopes for 7210 SAS-Dxp

On the 7210 SAS-Dxp, two slopes can be used per queue: a high-priority RED slope and a low-priority RED slope. The slope policy is only used for TCP traffic. Non-TCP traffic is always tail-dropped if the queues are full.

The high-priority RED slope manages access to the shared portion of the buffer pool for high-priority or in-profile TCP packets. The low-priority RED slope manages access to the shared portion of the buffer pool for low-priority or out-of-profile TCP packets. By default, the high-priority and low-priority RED slopes are disabled.

When 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, two RED slopes are used to determine buffer availability on a packet-by-packet basis. A TCP packet that is 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 the low-priority RED slope. Non-TCP packets are tail-dropped if the queue is full.

2.10.4.3 Simplified overview of RED for 7210 SAS-D and 7210 SAS-Dxp


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 used CBS (actually in use by queues, not just defined by the CBS) is oversubscribed and has stolen buffers from the shared size, lowering the effective shared buffer size equal to the shared buffer utilization size.
- If the packet is queued, a new shared buffer average utilization is calculated using the time average-factor (TAF) for the buffer pool. The TAF describes the weighting between the previous shared buffer average utilization result and the new shared buffer utilization in determining the new shared buffer average utilization. (See [Tuning the shared buffer utilization calculation on 7210 SAS-D and 7210 SAS-Dxp](#) for more information.)
- 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 describe 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.

2.10.4.4 Tuning the shared buffer utilization calculation on 7210 SAS-D and 7210 SAS-Dxp

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

Figure 4: Calculation for average shared buffer utilization

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

sw0533

where:

$SBAU_n$ = Shared buffer average utilization for event n

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

SBU = The instantaneous shared buffer utilization

TAF = The time average factor

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

Table 20: TAF impact on shared buffer average utilization calculation

| TAF | 2^{TAF} | Equates To | Shared buffer instantaneous utilization portion | Shared buffer average utilization portion |
|-----|-----------|------------|---|---|
| 0 | 2^0 | 1 | 1/1 (1) | 0 (0) |
| 1 | 2^1 | 2 | 1/2 (0.5) | 1/2 (0.5) |
| 2 | 2^2 | 4 | 1/4 (0.25) | 3/4 (0.75) |
| 3 | 2^3 | 8 | 1/8 (0.125) | 7/8 (0.875) |
| 4 | 2^4 | 16 | 1/16 (0.0625) | 15/16 (0.9375) |
| 5 | 2^5 | 32 | 1/32 (0.03125) | 31/32 (0.96875) |

| TAF | 2^{TAF} | Equates To | Shared buffer instantaneous utilization portion | Shared buffer average utilization portion |
|-----|------------------|------------|---|---|
| 6 | 2^6 | 64 | 1/64 (0.015625) | 63/64 (0.984375) |
| 7 | 2^7 | 128 | 1/128 (0.0078125) | 127/128 (0.9921875) |
| 8 | 2^8 | 256 | 1/256 (0.00390625) | 255/256 (0.99609375) |
| 9 | 2^9 | 512 | 1/512 (0.001953125) | 511/512 (0.998046875) |
| 10 | 2^{10} | 1024 | 1/1024 (0.0009765625) | 1023/2024 (0.9990234375) |
| 11 | 2^{11} | 2048 | 1/2048 (0.00048828125) | 2047/2048 (0.99951171875) |
| 12 | 2^{12} | 4096 | 1/4096 (0.000244140625) | 4095/4096 (0.999755859375) |
| 13 | 2^{13} | 8192 | 1/8192 (0.0001220703125) | 8191/8192 (0.9998779296875) |
| 14 | 2^{14} | 16384 | 1/16384 (0.00006103515625) | 16383/16384 (0.99993896484375) |
| 15 | 2^{15} | 32768 | 1/32768 (0.000030517578125) | 32767/32768 (0.999969482421875) |

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

2.10.4.5 Slope policy parameters for 7210 SAS-D

The elements required to define a slope policy are the following:

- a unique policy ID
- **the high and low RED slope shapes for the queues**
Settings for the high-priority and low-priority RED slopes or the high-slope (for TCP in-profile packets), low-slope (for TCP out-of-profile packets), and non-TCP slope (for non-TCP packets). All three slopes are on a per-port per-queue basis
- configurable parameters on each slope are start-avg, max-avg, max-prob and time averaging-factor (TAF)

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

Each access port egress buffer pool and access-uplink port egress buffer pool can be associated with only one slope policy ID.

The slope policy ID default is reserved for the default slope policy. The default policy cannot be deleted or changed. The default slope policy is applied to all access and network buffer pools that do not have another slope policy explicitly assigned.

The following table lists the default values for the default slope policy.

Table 21: Default slope policy definition for 7210 SAS-D

| Parameter | Description | Setting |
|---------------------|----------------------|------------------------------|
| Policy ID | Policy ID | Default (for default policy) |
| High (RED) slope | Administrative state | Shutdown |
| | start-avg | 70% utilization |
| | max-avg | 90% utilization |
| | max-prob | 75% probability |
| Low (RED) slope | Administrative state | Shutdown |
| | start-avg | 50% utilization |
| | max-avg | 75% utilization |
| | max-prob | 75% probability |
| Non-TCP (RED) slope | Administrative State | Shutdown |
| | start-avg | 50% utilization |
| | max-avg | 75% utilization |
| | max-prob | 75% probability |

2.10.4.6 Slope policy parameters for 7210 SAS-Dxp

The elements required to define a slope policy are the following:

- a unique policy ID
- **the high and low RED slope shapes for the queues**

Settings for the high-priority and low-priority RED slopes. Slope policies are only used for TCP traffic. Non-TCP traffic is always tail-dropped if the queues are full.

- configurable parameters on each slope are start-avg, max-avg, max-prob, and time averaging-factor (TAF)

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

Each access port egress buffer pool and access-uplink port egress buffer pool can be associated with one only slope policy ID. The slope policy ID default is reserved for the default slope policy. The default policy cannot be deleted or changed. The default slope policy is applied to all access and network buffer pools that do not have another slope policy explicitly assigned.

The following table lists the default values for the default slope policy.

Table 22: Default slope policy definition for 7210 SAS-Dxp

| Parameter | Description | Setting |
|------------------|----------------------|------------------------------|
| Policy ID | Policy ID | Default (for default policy) |
| High (RED) slope | Administrative state | Shutdown |
| | start-avg | 70% utilization |
| | max-avg | 90% utilization |
| | max-prob | 75% probability |
| Low (RED) slope | Administrative state | Shutdown |
| | start-avg | 50% utilization |
| | max-avg | 75% utilization |
| | max-prob | 75% probability |

2.11 QoS policy entities

Services are configured with default QoS policies. Additional policies must be explicitly created and associated. The following policies are configured by default:

- one default service ingress QoS policy
- one default service egress QoS policy
- one default access egress QoS policy
- one default network QoS policy
- one default port scheduler policy

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

When you create a new QoS policy, default values are provided for most parameters with the exception of the policy ID and 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**

On 7210 SAS-D and 7210 SAS-Dxp, SAP-ingress policies are supported on an Epipe.

- **VPLS**

On 7210 SAS-D and 7210 SAS-Dxp, SAP-ingress policies are supported on a VPLS SAP.

QoS policies can be applied to the following entities on 7210 SAS-D and 7210 SAS-Dxp:

- access-egress policies on access ports
- network QoS policy on access-uplink port
- network queue policy (egress) on access-uplink port

2.12 Configuration notes

The following information describes QoS implementation 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.



Note:

DEI classification and marking is only supported on the 7210 SAS-D and 7210 SAS-Dxp.

3.1 DEI-based classification

The DEI bit in the received packet is used to determine 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 profile assigned on ingress can be used to enable color-aware metering with SAP ingress policing and access-uplink port ingress policing. The profile of the packet can be reassigned by ingress meters/policers. When policing is used on SAP ingress, the final profile of the packet is determined by the meter/policers, based on the configured CIR/PIR rates. If a packet is below CIR rate, it is assigned in-profile/green and if it exceeds the CIR rate and is below the PIR rate, it is assigned out-of-profile/yellow.

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

The following support is available for DEI classification:

- Under the port configuration, a command is provided to enable DEI-based classification, which enables the use of DEI for ingress classification on a per-port basis. The initial profile (or color) is based on the DEI/CFI bit. 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 by the subsequent processing flow in the hardware. The FC classification can be done using MAC or IP criteria.
- All the SAPs configured on the port (access or hybrid) can use DEI classification for color-aware metering, if required. When DEI classification is enabled on the port, the user has an option to use color-blind metering for some SAPs and color-aware metering for other SAPs configured on the same port. When using color-blind mode, the ingress profile assigned to the packet based on the DEI bit is ignored.
- The user is provided with an option in the SAP-ingress policy to configure a policer as color aware or color blind. In color-aware mode, the DEI bit in the packet determines the ingress profile of the packet. If the user configures the meter/policer mode as color-aware, the DEI bit of the incoming packet is used by the policer as the ingress profile.
- When using policing, the final profile of the packet is assigned by the ingress meter (based on configured CIR/PIR rate) in both color-aware and color-blind mode.
- For network port policies, DEI-based classification is supported only when dot1p classification criteria are in use; it cannot be used when DSCP-based classification is used.

3.2 DEI-based marking

DEI-based marking is supported on access ports, access-uplink ports, and network ports. 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 network) to the subsequent nodes in the network. It allows high-priority in-profile packet to be allocated appropriate resources by all the network nodes on the path to the final destination. Similarly, it allows out-of-profile packets to be treated with less preference compared to in-profile packets by all the network nodes on the path to the final destination. The egress marking behavior must be symmetric to the ingress classification behavior.

The following support is available for DEI-based marking:

- option to mark DEI bits for port egress on access ports and access-uplink ports
- 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 an 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.



Note:

See [Network QoS policy command reference](#) and [Access egress QoS policy command reference](#), and see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for information about the CLI commands for DEI.

3.3 Configuration guidelines

The following are configuration guidelines for DEI-based classification and marking:

- While disabling DEI-based classification on a port, all the meters used by the SAPs configured on this port must be in color-blind mode. The converse is also true, that is, while attaching a SAP-ingress QoS policy with meter as color aware to a SAP, the DEI-based classification must be enabled on the port where the SAP exists.
- While configuring DEI-based classification in a access-uplink network QoS Policy (ingress), only dot1p classification can be used.
- DEI classification must be disabled on that port before changing the mode (access, access-uplink, network, hybrid) to another mode.
- All the ports under a LAG should have the same configuration for DEI classification (either enabled or disabled for all member ports of a LAG). If the LAG configuration changes, the port configuration is also updated accordingly. Port configuration under the LAG cannot be changed.

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.

4.1.1 Applications

This feature is useful when connecting the 7210 SAS to an Ethernet-over-SDH (EoSDH) (or microwave) network, where the network allocates predetermined bandwidth to the nodes connecting into it, based on the transport bandwidth requirement. When connecting to such a network, it is important that the traffic sent into the SDH node does not exceed the configured values, because the SDH network does not have the QoS capabilities and buffers required to prioritize the ingress traffic.

Egress rate attributes include the following:

- Allows for per port configuration of the maximum egress port rate, 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 7210 SAS-D and 7210 SAS-Dxp provide support for a burst parameter to control the amount of burst the egress port can generate.
- When ports are members of a LAG, all the ports use the same value for **egress-rate** and **max-burst**.
- If frame overhead accounting (also known as Frame-based accounting) is enabled, the queue scheduler accounts for the Ethernet frame overhead.

4.1.2 Effect of port-level rate-limiting on queue functionality

When an **egress-rate sub-rate** value is provided, the access-uplink port egress queue rates that are specified using percentages use the **egress-rate** value instead of the port bandwidth (if egress rate is less than port bandwidth) to configure the queue rates. Configuring the egress port rate to different values results in a corresponding dynamic adjustment of rates for the egress queues configured on the access-uplink ports.

When the **egress-rate sub-rate** value is set, the CBS and MBS of the associated network queues do not change.

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).
- For 7210 SAS-D devices, the **max-burst** command configures a maximum burst (in kilobits) associated with the egress-rate. This parameter is optional and if it is not defined, by default it is set to 32 kb for a 1G port. Users cannot configure **max-burst** without configuring **egress-rate**. The value should be between 32 and 16384 or default.
- For 7210 SAS-Dxp devices, the **max-burst** command configures a maximum-burst (in kilobits) associated with the egress-rate. This parameter is optional and if it is not defined, by default it is set to 65 kb for a 1G port and 98 kb for a 10G port. Users cannot configure **max-burst** without configuring **egress-rate**. The value should be between 32 and 16384 or default.
- By default the **egress-rate** command is not set on the port. The default **egress-rate** value for a port is the maximum (which is equal to the line rate).

Example

The following sample configuration output shows the **egress-rate** configuration for a port.

```
*A:Dut-1>config>port# info
-----
    ethernet
        egress-rate 120000 max-burst 234
    exit
    no shutdown
-----
*A:Dut-1>config>port#
```

4.2.1 Modifying port level egress-rate command

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

Example

The following sample configuration shows a modified **egress-rate** configuration for a port.

```
*A:Dut-1>config>port# ethernet egress-rate 10000 max-burst default
*A:Dut-1>config>port# info
-----
    ethernet
        egress-rate 10000
    exit
    no shutdown
-----
*A:Dut-1>config>port#
```

4.2.2 Removing port level egress-rate command

To remove 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 sample configuration output 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#
```

4.2.2.1 Default egress-rate values

By default, no **egress-rate** is configured for a port. For more information about the CLI and description, see [Port level egress-rate command reference](#).

4.3 Port level egress-rate command reference

4.3.1 Command hierarchies

4.3.1.1 Configuration commands

```
- config
  - port
    - ethernet
      - egress-rate sub-rate [max-burst size-in-kbits]
      - no egress-rate
```

4.3.1.2 Show commands

```
- show
  - port [port-id]
```

4.3.2 Configuration descriptions

4.3.2.1 Configuration commands

egress-rate

Syntax

egress-rate <sub-rate> [**max-burst** <size-in-kbits>]
no egress-rate

Context

config>port>ethernet

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures maximum egress rate and corresponding burst value for a port. The egress rate is configured as kb/s while max-burst is configured as kilobits while max-burst should be between 64 and 16384 or default.



Note:

10G ports are not supported on the 7210 SAS-D.

The **no** form of the command removes the egress rate from the port.

Default

No egress-rate and max-burst is configured for the port.

Parameters

sub-rate

Specifies an integer value between 1 and 1000000 kb/s and between 1 and 10000000 kb/s for 10G port. 7210 SAS-D devices do not support 10G port.

max-burst size-in-kbits

Specifies an integer value, in kilo-bits, between 32 Kbits and 16384 Kbits the default value is 64 Kbits.

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 the 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 23: Output fields: show PoE port \(Ethernet\)](#) 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
Link-level       : Ethernet
Admin State      : up
Oper State       : up
Physical Link     : Yes
IfIndex          : 35684352
Last State Change : 01/17/2011 04:05:37
Last Cleared Time  : N/A

Oper Speed       : 1 Gbps
Config Speed     : 1 Gbps
Oper Duplex      : full
Config Duplex    : full
MTU              : 1514
Hold time up     : 0 seconds
Hold time down   : 0 seconds

Configured Mode   : access
Dot1Q Ethertype   : 0x8100
Net. Egr. Queue Pol : default
Egr. Sched. Pol   : default
Auto-negotiate    : limited
Accounting Policy  : None
Egress Rate       : Default
Uplink            : No

Encap Type        : null
QinQ Ethertype    : 0x8100
Access Egr. Qos *: 1
Network Qos Pol   : n/a
MDI/MDX           : MDI
Collect-stats     : Disabled
Max Burst         : Default

Down-when-looped  : Disabled
Loop Detected     : False

Keep-alive        : 10
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     : 3HE00027AAAA02 ALA IPU1AELDAB=
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
=====
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
=====

```

Table 23: Output fields: show PoE port (Ethernet)

| 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 |

| Label | Description |
|-------------------|--|
| 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 |
| 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 will not use any tags or labels to delineate a service dot1q — ingress frames carry 802.1Q tags, where each tag signifies a different service |

| Label | Description |
|---------------------|--|
| | 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. |
| 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 |

| Label | Description |
|---------------------------|---|
| Retry | The minimum wait time before the port is re-enabled after it is brought down because of a loop detection |
| Configured Address | The base chassis Ethernet MAC address |
| Hardware Address | The interface hardware- 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 |
| 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. |

| Label | Description |
|-------------------------------------|---|
| Errors Input/Output | <p>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.</p> <p>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.</p> |
| 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 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 |

| Label | Description |
|--|--|
| | 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 because of excessive collisions |
| Too long Frames | The number of frames received that exceed the maximum permitted frame size |
| Int MAC Tx Errs | The number of frames for which a transmission fails because 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 because of an internal MAC sublayer receive error |

5 SAP egress aggregate meter

This chapter provides information to configure service-level egress rate limits using the command line interface.

5.1 Overview

The SAP egress aggregated meter feature allows a user to limit the amount of traffic sent out of a SAP by enforcing service-level egress rate limits.

This feature is supported only on the 7210 SAS-D and 7210 SAS-Dxp.

This feature allows a meter to be associated with the SAP egress to limit the aggregate amount of traffic across all FCs that are sent out of the SAP. This feature is supported for SAPs in P2MP services (Epipe, VPLS, RVPLS, IES, and VPRN).

5.1.1 Configuration notes

- Before configuring SAP egress aggregate meter on a SAP, the user must reallocate the resources from the egress-internal-tcam pool toward the SAP egress meters using the CLI command **configure>system>resource-profile>egress-internal-tcam>egress-sap-aggregate-meter**. For more on the **egress-sap-aggregate-meter** command, see the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide. The egress-internal- tcam pool resources are shared with other features such as egress ACLs.
- The SAP egress aggregate policer is not FC and profile aware. The packets scheduled out of the per-port queues are on a first-come-first-served basis and limits the amount of traffic to the configured rate across all the FCs for a SAP.
- For more information about the optional parameter **enable-stats** of the **configure>service>vpls>sap>egress>aggregate-meter-rate** CLI command, see the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide. This option associates a counter with the meter to count the forwarded packets/octets and dropped packets/octets. If the option and accounting is enabled for the SAP, service egress accounting records contain the forwarded and dropped counts.

5.2 Basic configurations

- Allocate resources to the feature using the command **configure>system> resource-profile>egress-internal-tcam>egress-sap-aggregate-meter**. If necessary, free up resources for use by this feature by taking away resources from other features.
- Configure the SAP egress aggregate meter rate using the command **configure>service>vpls>sap>egress>aggregate-meter-rate rate-in-kbps[burst burst-in-kbits] [enable-stats]**



Note:

Use the **enable-stats** parameter to enable counter used to count total forwarded packet out of the SAP.

5.3 SAP egress aggregate meter command reference

For more information of the SAP egress aggregate meter commands, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide*.

6 Frame-based accounting

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

6.1 Overview

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 20 bytes (12 + 8). The QoS meter/policer and shaper use this overhead when allocating bandwidth for Ethernet ports.

6.1.1 Frame-based accounting

A configurable CLI command enables accounting of the frame overhead on ingress or egress. This is a system-wide parameter that affects the behavior of the ingress meter or egress rate. When disabled, the queue rates and egress-rate do not account for the Ethernet frame overhead. By default, frame-based accounting is disabled for both ingress and egress.



Note:

Frame-based accounting for SAP-egress aggregate meters is not supported on the 7210 SAS platforms described in this document.

6.1.2 Effects of enabling ingress frame-based accounting on ingress meter functionality

To enable system-wide consistency in configuring QoS queue and meter rate parameters, the meters used on the system ingress may need to account for Ethernet frame overhead. Access uplink ingress and service ingress meters account for Ethernet frame overhead. A configurable CLI command can enable or disable the frame overhead accounting. This is a system-wide command that affects the behavior of all the meters in the system.

6.1.3 Effects of enabling egress frame-based accounting on access-uplink queue functionality

If frame overhead consideration is enabled, the queue scheduler accounts for the Ethernet frame overhead. The maximum egress bandwidth accounts for the Ethernet frame overhead (it accounts for the IFG (inter-frame gap) and the preamble). Typically, the IFG and preamble constitute about 20 bytes (12 + 8). The overhead for Ethernet ports uses this value.

A configurable CLI command enables accounting of the frame overhead. This is a system-wide command that affects the behavior of all egress queues (when frame-based accounting is enabled on egress port, the associated queues also account for frame overhead implicitly). When disabled, the port egress-rate command does not account for the Ethernet frame overhead.



Note:

Frame-based accounting does not affect the SAP egress aggregate rate command on 7210 SAS-D and 7210 SAS-Dxp. In other words, the SAP egress aggregate command does not account for the Ethernet frame overhead regardless of whether egress frame-based accounting is enabled or disabled.

6.1.4 Accounting and statistics

Accounting records and statistics do not account for frame overhead.

6.2 Basic configurations

To enable frame-based accounting, perform the following:

- Enable the **frame-based-accounting** command in the **config>qos>frame-based-accounting** context.
- Configure the **ingress-enable** command to enable frame-based accounting for ingress metering.
- Configure the **egress-enable** command to enable frame-based-accounting for egress queue rates, scheduler and port level egress-rate.

Example

The following is a sample frame-based accounting configuration output.

```
*A:Dut-1>config>qos>frame-based-accounting# info detail
-----
no ingress-enable
no egress-enable
-----
*A:Dut-1>config>qos>frame-based-accounting#
```

6.2.1 Enabling and disabling frame-based accounting

To enable the **frame-based-accounting** command for ingress, use the **ingress-enable** command. To disable the **frame-based-accounting** command for ingress, use the **no ingress-enable** command. To enable the **frame-based-accounting** command on egress, use the **egress-enable** command. To disable **frame-based-accounting** command on egress, use the **no egress-enable** command.

```
config>qos>frame-based-accounting
```

Example

The following is a sample configuration output that shows the enabling of frame-based-accounting.

```
*A:Dut-1>config>qos>frame-based-accounting# ingress-enable
*A:Dut-1>config>qos>frame-based-accounting# egress-enable
*A:Dut-1>config>qos>frame-based-accounting# info
-----
ingress-enable
egress-enable
-----
```

```
*A:Dut-1>config>qos>frame-based-accounting#
```

Example

The following is a sample configuration output that shows the disabling of frame-based-accounting.

```
*A:Dut-1>config>qos>frame-based-accounting# no ingress-enable
*A:Dut-1>config>qos>frame-based-accounting# no egress-enable
*A:Dut-1>config>qos>frame-based-accounting# info detail
-----
                no ingress-enable
                no egress-enable
-----
*A:Dut-1>config>qos>frame-based-accounting#
```

6.3 Frame-based accounting command reference

6.3.1 Command hierarchies

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

6.3.1.1 Configuration commands

```
- config
  - qos
    - frame-based-accounting
      - [no] egress-enable
      - [no] ingress-enable
```

6.3.1.2 Show commands

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

6.3.2 Configuration descriptions

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

6.3.2.1 Configuration commands

frame-based-accounting

Syntax

frame-based-accounting

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure frame-based accounting.

egress-enable

Syntax

[no] egress-enable

Context

config>qos>frame-based-accounting

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the frame-based accounting for access egress, network queue, port scheduler, SAP or network aggregate rate, and port-level egress rate.

The **no** form of the command disables frame-based accounting for all egress QoS.

Default

no egress-enable

ingress-enable

Syntax

[no] ingress-enable

Context

config>qos>frame-based-accounting

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables frame-based accounting for SAP-ingress and network QoS.

The **no** form of the command disables frame-based accounting for SAP-ingress and network QoS.

Default

no ingress-enable

6.3.2.2 Show commands

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 the 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 it is shown as packet-based.

Parameters

policy-id

Displays information about the specific policy ID.

association

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 SAP-ingress QoS policy information, and [Table 24: Output fields: QoS SAP-ingress](#) describes the output fields.

Sample output

```
*A:Dut-1# show qos sap-ingress 1
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (1)
-----
Policy-id      : 1                Scope      : Template
Default FC     : be
Criteria-type  : None
Accounting     : frame-based
Classifiers Allowed: 16          Meters Allowed : 8
Classifiers Used : 2             Meters Used   : 2
Description    : Default SAP ingress QoS policy.
=====
*A:Dut-1#
```

Table 24: Output fields: QoS SAP-ingress

| Label | Description |
|---------------|--|
| Policy-Id | Displays 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 FC 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 FC for matched traffic. |
| Accounting | Packet-based Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow. |
| | Frame-based Specifies that the meters associated with this policy account for the packet framing overheads (such as for Ethernet the IFG and |

| Label | Description |
|---------------------|---|
| | preamble), while accounting the bandwidth to be used by the flow. |
| Classifiers Allowed | Indicates the number of classifiers allowed for a service. |
| Meters Allowed | Indicates the number of meters allowed for a service. |
| Classifiers Used | Indicates the number of classifiers used for a service. |
| Meters Used | Indicates the number of meters used for a service. |
| Description | A text string that helps identify the policy's context in the configuration file. |

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 in addition to other details about the policy. When **frame-based-accounting** is enabled, accounting is shown as frame-based, otherwise it is shown as packet-based.

Parameters

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

Output

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

Sample output

```
*A:Dut-1# show qos network 1
=====
QoS Network Policy
=====
```

```

-----
Network Policy (1)
-----
Policy-id      : 1                      Remark      : False
Forward Class  : be                    Profile      : Out
Attach Mode    : l2                    Config Mode  : l2+mpls
Scope          : Template              Policy Type  : port
Accounting     : frame-based
Description    : Default network-port QoS policy.
-----
Meter Mode    CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS    MBS
-----
1      TrTcm_CA  0          closest    max      closest  32     128
-----
FC                      UCastM      MCastM
-----
No FC-Map Entries Found.
-----
Port Attachments
-----
Port-id : 1/1/3
Port-id : 1/1/6
Port-id : 1/1/7
Port-id : 1/1/8
Port-id : 1/1/9
=====
*A:Dut-1#

```

Table 25: Output fields: show QoS network

| Label | Description |
|---------------|---|
| Policy-Id | The ID that uniquely identifies the policy. |
| Remark | Remarking can be enabled or disabled. |
| Forward Class | Specifies the FC name. |
| Profile | Out Specifies the EXP marking for the packets which are out-of-profile, egressing on this queue. Specifies the dot1p marking for the out-of-profile packets which egressing on this queue. |
| | In Specifies the EXP marking for the packets which are in-of-profile, egressing on this queue. Specifies the dot1p markings for in-profile packets egressing on this queue. |
| Scope | Exclusive Specifies that this policy can be applied only to a single network port. |
| | Template Specifies that this policy can be applied to multiple network ports on the router. |

| Label | Description |
|-------------|--|
| Policy Type | Specifies the policy type. |
| Accounting | <p>Packet-based</p> <p>Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow.</p> <p>Frame-based</p> <p>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.</p> |
| Description | A text string that helps identify the policy's context in the configuration file. |
| Meter Mode | Specifies the configured mode of the meter. |
| 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 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 |

| Label | Description |
|---------|---|
| | The operational PIR for the meters will be the rate closest to the rate specified using the rate command without exceeding the operational PIR. |
| 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. |
| Port-id | Specifies the port number. |

access-egress

Syntax

access-egress [*policy-id*] [**association** | **detail**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the accounting status of an access-egress policy in addition to other details about the policy. When **frame-based-accounting** is enabled, accounting is shown as frame-based, otherwise it is shown as packet-based.

Parameters

policy-id

Displays information about the specific policy ID.

association

Displays the policy associations.

detail

Displays the policy information in detail.

Output

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

Sample output

```
*A:Dut-1# show qos access-egress 1
=====
QoS Access Egress
=====
-----
Policy-id      : 1                      Scope      : Template
Remark        : False
Accounting     : frame-based
Description    : Default Access egress QoS policy.
=====
*A:Dut-1#
```

Table 26: Output fields: access egress

| Label | Description |
|-------------|---|
| Policy-Id | Specifies the ID that uniquely identifies the policy. |
| Scope | Exclusive Specifies that this policy can be applied only to a single access egress port. Template Specifies that this policy can be applied to multiple access ports on the router. |
| Remark | True Specifies that 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 Specifies that remarking is disabled for the policy. |
| Accounting | Packet-based Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow. Frame-based Specifies that the meters associated with this policy account for the packet framing overheads (such as for Ethernet the IFG and preamble), while accounting the bandwidth to be used by the flow. |
| Description | A text string that helps identify the policy’s context in the configuration file |

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 the accounting status of a network-queue policy in addition to other details about 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 network queue information, and [Table 27: 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
=====
*A:Dut-1#
```


Table 27: Output fields: network queue

| Label | Description |
|-------------|--|
| Policy | The policy name that uniquely identifies the policy. |
| Accounting | <p>Packet-based</p> <p>Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow.</p> <p>Frame-based</p> <p>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.</p> |
| 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. |

port-scheduler-policy

Syntax

port-scheduler-policy [*port-scheduler-policy-name*] [**association**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the accounting status of a port-scheduler policy in addition to other details about the policy. When **frame-based-accounting** is enabled, accounting is shown as frame-based, otherwise it is shown as packet-based.

Parameters

port-scheduler-policy-name

Displays information about the specific port scheduler policy.

association

Displays the associations of the port scheduler policy.

Output

The following output is an example of port scheduler policy information, and [Table 28: Output fields: port scheduler policy](#) describes the output fields.

Sample output

```
*A:Dut-1# show qos port-scheduler-policy default
=====
QoS Port Scheduler Policy
=====
Policy-Name      : default
Description      : Default Port Scheduler policy.
Accounting       : frame-based
Mode             : STRICT
Last changed     : 08/06/2001 18:36:04
Number Of Queues : 8
```

Table 28: Output fields: port scheduler policy

| Label | Description |
|--------------|---|
| Policy-Name | Displays the port scheduler policy name. |
| Mode | Displays the port scheduler policy mode (STRICT, RR, WRR, WDRR). |
| Accounting | Packet-based Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow. Frame-based Specifies that the meters associated with this policy account for the packet framing overheads (such as for Ethernet the IFG and preamble), while accounting the bandwidth to be used by the flow. |
| Last Changed | Displays the last time the configuration changed. |
| Queue # | Displays the weight of the queue if configured. |

7 Network QoS policies

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

7.1 Overview of network QoS policies

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

The ingress component of the port policy defines how the dot1p bits are mapped to the internal FC and profile state. The FC and profile state define the PHB and the QoS handling 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 meters for unicast traffic and multipoint traffic.



Note:

7210 SAS platforms provide different mechanisms to limit the bandwidth per FC. On 7210 SAS-D and 7210 SAS-Dxp, the user needs to use policers/meters to rate-limit the traffic per FC.

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

Access uplink port egress marking support consists of the following:

- For packets sent out of a 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.
- On 7210 SAS-D and 7210 SAS-Dxp, remarking can be enabled or disabled on access uplink ports. An option is provided to map FC to dot1p and DEI bits and IP DSCP values.

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 "CLI Usage" in the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide for information about the tasks and commands necessary to access the command line interface, and to configure and maintain your devices.

7.1.1 Resource allocation for network QoS policy

This section describes the allocation of QoS resources for network QoS policies. When the port mode is changed to access-uplink, a default network QoS policy is applied. For the default policy, two meters and eighteen classification entries in hardware are allocated.



Note:

The number of resources used per network QoS policy determines the number of access-uplink ports that can be configured. If more resources are used, fewer access-uplink ports can be used, and the other way around.

For every FC in use, the system allocates two classification entries in hardware, provided the FC is configured to use both the unicast meter and the multicast meter, or provided the default meter 9 is configured in the policy. If multiple match criteria entries map to the same FC, each of these are allocated two classification entries in hardware. For example, if there are two match-criteria entries that map to FC "af" a total of four classification entries are allocated in hardware, and if there are four match-criteria entries that map to FC "af" a total of eight classification entries are allocated in hardware.

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

To calculate the number of QoS resources used by an access-uplink port, determine the following items:

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

Only the FCs used by the match-criteria classification entries are considered for the "number of FCs" and are therefore referred to as "FCs in use". Also, in network policies of the type **ip-interface**, the default multipoint meter 9 is created in a policy; whereas, for policy of type "port", the default multipoint meter needs to be explicitly configured by the user, if required.

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

- If an FC is in use and is created without explicit meters, use default meter 1 for unicast traffic and default meter 9 (if configured) for all other traffic types (that is, broadcast, multicast, and unknown-unicast). This requires two classification entries in hardware. If default multipoint meter 9 is not configured, the FC uses the unicast meter for all traffic types. In this case, the FC requires a single classification entry in hardware.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter 9 (if configured) for all other traffic types. This requires two classification entries in hardware. If default multipoint meter 9 is not configured, the FC uses the unicast meter for all traffic types. In this case, the FC requires a single classification entry in hardware.
- If an FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic. This requires two classification entries in hardware.

Given the number of match criteria and the number of FCs in use, use the following equation to calculate the total number of classification entries per policy (for example TC):

- $TC = 2 * E(i)$
- $i = nc, h1, ef, h2, l1, af, l2, be$

where:

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



Note:

In the worst case, only 2 classification entries are used per FC in a network policy, because only two traffic types are supported.

Determine number of policers or meters to use (for example TP). A maximum of 16 meters per network policy are available.

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

7.1.1.1 Network QoS policies resource usage examples

This section provides usage examples for network QoS policies.

7.1.1.1.1 Example 1

```
network 1 create
  description "default QoS policy"
  ingress
    default-action fc be profile out
    meter 1 create
    exit
    meter 9 multipoint create
    exit
  exit
  egress
    fc af
    exit
    fc be
    exit
    fc ef
    exit
    fc h1
    exit
    fc h2
    exit
    fc l1
    exit
    fc l2
    exit
    fc nc
    exit
  exit
```

The number of classification entries (TC) used is calculated as follows.

$$(2 * 0)nc + (2 * 0)h1 + (2 * 0)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 0)af + (2 * 0)l2 + (2 * 1)be = 18$$

The number of meters (TP) used is 2 (meter 1 and 9).

7.1.1.1.2 Example 2

```
network 2 create
  description "network-policy-2"

  ingress
    default-action fc be profile out
    meter 1 create
    exit
    meter 2 create
    exit
    meter 9 multipoint create
    exit
    meter 12 multipoint create
    exit
    fc "af" create
      meter 2
      multicast-meter 12
    exit
    dot1p 2 fc af profile out
  exit
  egress
    fc af
    exit
    fc be
    exit
    fc ef
    exit
    fc h1
    exit
    fc h2
    exit
    fc l1
    exit
    fc l2
    exit
    fc nc
    exit
  exit
exit
```

The number of classification entries (TC) used is calculated as follows.

$$(2 * 0)nc + (2 * 0)h1 + (2 * 0)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 1)af + (2 * 0)l2 + (2 * 1)be = 4$$

The number of meters (TP) used is 4 (meters 1,2,9,12).

7.1.1.1.3 Example 3

```
network 3 create
  description "network-policy-3"
  ingress
    default-action fc be profile out
    meter 1 create
    exit
    meter 2 create
    exit
```

```

        meter 9 multipoint create
        exit
        meter 12 multipoint create
        exit
        fc "af" create
            meter 2
            multicast-meter 12
        exit
        fc "be" create
            meter 2
            multicast-meter 12
        exit
        dotlp 2 fc af profile out
    exit
    egress
        fc af
        exit
        fc be
        exit
        fc ef
        exit
        fc h1
        exit
        fc h2
        exit
        fc l1
        exit
        fc l2
        exit
        fc nc
        exit
    exit
exit

```

The number of classification entries (TC) used is calculated as follows.

$$(2 * 0)nc + (2 * 0)h1 + (2 * 0)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 1)af + (2 * 0)l2 + (2 * 1)be = 4$$

The number of meters (TP) used is 2 (meters 2 and 12).

7.1.1.1.4 Example 4

```

network 4 create
    description "network-policy-4"
    ingress
        default-action fc be profile out
        meter 1 create
        exit
        meter 9 multipoint create
        exit
        dotlp 1 fc l2 profile in
        dotlp 2 fc af profile out
        dotlp 3 fc af profile in
        dotlp 4 fc h2 profile in
        dotlp 5 fc ef profile in
        dotlp 6 fc h1 profile in
        dotlp 7 fc nc profile in
    exit
    egress

```

```

        fc af
        exit
        fc be
        exit
        fc ef
        exit
        fc h1
        exit
        fc h2
        exit
        fc l1
        exit
        fc l2
        exit
        fc nc
        exit
    exit
exit

```

The number of filter entries (TC) used is calculated as follows.

$$(2 * 1)nc + (2 * 1)h1 + (2 * 1)ef + (2 * 1)h2 + (2 * 0)l1 + (2 * 2)af + (2 * 1)l2 + (2 * 1)be = 16$$

The number of meters (TP) used is 2 (meters 1 and 9).

7.1.1.1.5 Example 5

```

network 5 create
  description "network-policy-5"
  ingress
    default-action fc be profile out
    meter 1 create
    exit
    meter 2 create
    exit
    meter 9 multipoint create
    exit
    meter 12 multipoint create
    exit
    fc "af" create
    exit
    fc "be" create
    exit
    fc "ef" create
    exit
    fc "h1" create
    exit
    fc "h2" create
    exit
    fc "l2" create
    exit
    fc "nc" create
    exit
    dot1p 1 fc l2 profile in
    dot1p 2 fc af profile out
    dot1p 3 fc af profile in
    dot1p 4 fc h2 profile in
    dot1p 5 fc ef profile in
    dot1p 6 fc h1 profile in

```



```

        dot1p 7 fc nc profile in
    exit
    egress
        fc af
        exit
        fc be
        exit
        fc ef
        exit
        fc h1
        exit
        fc h2
        exit
        fc l1
        exit
        fc l2
        exit
        fc nc
        exit
    exit

```

The number of classification entries (TC) used is calculated as follows.

$$(2 * 1)nc + (2 * 1)h1 + (2 * 1)ef + (2 * 1)h2 + (2 * 0)l1 + (2 * 2)af + (2 * 1)l2 + (2 * 1)be = 16$$

The number of meters (TP) used is 2 (meters 1 and 9). Note that meters 2 and 12 are not accounted for, because they are not associated with any FC.

7.1.1.1.6 Example 6

```

network 6 create
    description "network-policy-6"

    ingress
        default-action fc be profile out
        meter 1 create
        exit
        meter 2 create
        exit
        meter 3 create
        exit
        meter 9 multipoint create
        exit
        meter 12 multipoint create
        exit
        fc "af" create
            meter 2
            multicast-meter 12
        exit
        fc "be" create
        exit
        fc "ef" create
        exit
        fc "h1" create
            meter 3
        exit
        fc "h2" create
        exit
        fc "l2" create
        exit

```

```

        fc "nc" create
        meter 3
    exit
    dot1p 1 fc l2 profile in
    dot1p 2 fc af profile out
    dot1p 3 fc af profile in
    dot1p 4 fc h2 profile in
    dot1p 5 fc ef profile in
    dot1p 6 fc h1 profile in
    dot1p 7 fc nc profile in
exit
egress
    fc af
    exit
    fc be
    exit
    fc ef
    exit
    fc h1
    exit
    fc h2
    exit
    fc l1
    exit
    fc l2
    exit
    fc nc
    exit
exit
exit

```

The number of classification entries (TC) used is calculated as follows.

$$(2 * 1)nc + (2 * 1)h1 + (2 * 1)ef + (2 * 1)h2 + (2 * 0)l1 + (2 * 2)af + (2 * 1)l2 + (2 * 1)be = 16$$

The number of meters (TP) used is 5 (meters 1,2,3,9, and 12).

7.1.1.1.7 Example 7

```

network 2 create
description "network-policy 2"
scope template
ingress
    default-action fc be profile out
    meter 1 create
        mode trtcm
        adaptation-rule cir closest pir closest
        rate cir 0 pir max
        mbs default
        cbs default
    exit
    meter 9 multipoint create
        mode trtcm
        adaptation-rule cir closest pir closest
        rate cir 0 pir max
        mbs default
        cbs default
    exit
network-policy 2 0 fc be profile out

```

```

network-policy 2 1 fc l2 profile in
network-policy 2 2 fc af profile out
network-policy 2 3 fc af profile in
network-policy 2 4 fc h2 profile in
network-policy 2 5 fc ef profile in
network-policy 2 6 fc h1 profile in
network-policy 2 7 fc nc profile in
exit
egress
no remarking

```

The number of classification entries (TC) used is 18.

The number of meters (TP) used is 2.

7.1.1.1.8 Example 8

```

network 8 create
description "network-policy-8"
ingress
  default-action fc nc profile in
  meter 1 create
  exit
  meter 2 create
  exit
  meter 3 create
  exit
  meter 4 create
  exit
  meter 5 create
  exit
  meter 7 multipoint create
  exit
  meter 8 multipoint create
  exit
  meter 9 multipoint create
  exit
  meter 12 multipoint create
  exit
  fc "af" create
    meter 2
    multicast-meter 12
  exit
  fc "ef" create
    meter 4
    multicast-meter 8
  exit
  fc "h2" create
  exit
  fc "l2" create
    meter 3
    multicast-meter 7
  exit
  fc "nc" create
    meter 4
    multicast-meter 8
  exit
  dotlp 1 fc l2 profile in
  dotlp 3 fc af profile in
  dotlp 5 fc ef profile in
  dotlp 7 fc nc profile in

```

```
exit
egress
```

The number of classification entries (TC) used is calculated as follows.

$$(2 * 2)nc + (2 * 0)h1 + (2 * 1)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 1)af + (2 * 1)l2 + (0 * 0)be = 10$$

The numbers of meters (TP) used is 6 (meters 2, 3, 4, 7, 8, and 12).

7.1.2 Basic configuration

A basic network QoS policy must:

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

7.1.3 Create a network QoS policy

Configuring and applying QoS policies other than the default policy is optional. A default network policy of the appropriate type is applied to each 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) map. Otherwise, the default values are applied:
 - **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. Remarking can be enabled or disabled.
 - **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-profile or out-of-profile.
- Specify ingress criteria using the following:
 - **default action**
Defines the default action to be taken for packets that have an undefined dot1p bit set. The default-action specifies the FC to which such packets are assigned.

– dot1p

This specifies the dot1p 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.

```
config>qos#
network policy-id [network-policy-type network-policy-type]
description description-string
scope {exclusive|template}
egress
    remarking

    fc {be|l2|af|l1|h2|ef|h1|nc}
        dot1p-in-profile dot1p-priority
        dot1p-out-profile dot1p-priority
    default-action fc {fc-name} profile {in|out}
    dot1p dot1p-priority fc {fc-name} profile {in|out}
    fc {fc-name}
        meter {meter-id}
        multicast-meter {id}
    meter meter-id [multipoint]
        adaptation-rule cir {closest | max | min} pir {closest | max | min}
        cbs {size-in-kbits}
        mbs {size-in-kbits}
        mode {trtcm | srtcm}
        rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
```

The following commands associate a network QoS policy with the access-uplink port.

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

Example

The following is a sample configuration output for 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#
```

7.1.4 Default network policy values

The default network policy access uplink port is identified as policy-id "1". Default policies cannot be modified or deleted. The following table lists default network policy parameters.

Table 29: Default network policy 1

| Field | Default |
|-----------------|----------------------------|
| description | Default network QoS policy |
| Ingress | |
| default-action | fc be profile out |
| meter | 1 |
| mode | trtcm1 |
| adaptation-rule | cir closest |
| | pir closest |
| rate | cir 0 |
| | pir max |
| mbs | default kbits |
| cbs | default kbits |

7.1.5 DSCP marking for CPU-generated traffic

DSCP marking for CPU-generated traffic is not configurable by the user. The following table lists the default values.



Note:

DSCP and dot1p values in the table are applicable when remarking is disabled at the port level.

Table 30: DSCP and dot1p marking

| 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 |

| Protocol | IPv4 | DSCP marking | dot1p marking | Default FC | DSCP values | dot1p values |
|-------------------|------|--------------|---------------|-------------------|-------------------|--------------|
| 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 | NC | 48 | 7 |
| PTP ¹² | Yes | Yes | Yes | See ¹² | See ¹² | 7 |
| RADIUS | Yes | Yes | Yes | H2 | 34 | 4 |
| SCP | Yes | Yes | 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 | 4 |
| 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 |

7.1.6 Default DSCP mapping table

The following table lists default DSCP mapping values.

Table 31: Default DSCP mapping table

| DSCP name | DSCP value decimal | DSCP value hexadecimal | DSCP value binary | Label |
|-----------|--------------------|------------------------|-------------------|-------|
| Default | 0 | 0x00 | 0b000000 | be |

¹² 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.

| DSCP name | DSCP value decimal | DSCP value hexadecimal | DSCP value binary | Label |
|-----------------------|--------------------|------------------------|-------------------|---------|
| nc1 | 48 | 0x30 | 0b110000 | h1 |
| nc2 | 56 | 0x38 | 0b111000 | nc |
| ef | 46 | 0x2e | 0b101110 | ef |
| af11 | 10 | 0x0a | 0b001010 | assured |
| af12 | 12 | 0x0c | 0b001100 | assured |
| af13 | 14 | 0x0e | 0b001110 | assured |
| af21 | 18 | 0x12 | 0b010010 | l1 |
| af22 | 20 | 0x14 | 0b010100 | l1 |
| af23 | 22 | 0x16 | 0b010110 | l1 |
| af31 | 26 | 0x1a | 0b011010 | l1 |
| af32 | 28 | 0x1c | 0b011100 | l1 |
| af33 | 30 | 0x1d | 0b011110 | l1 |
| af41 | 34 | 0x22 | 0b100010 | h2 |
| af42 | 36 | 0x24 | 0b100100 | h2 |
| af43 | 38 | 0x26 | 0b100110 | h2 |
| default ¹³ | 0 | | | |

7.2 Service management tasks

This section provides information about service management tasks.

7.2.1 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.

¹³ The default FC mapping is used for all DSCP names or values for which there is no explicit FC mapping.

7.2.2 Remove a policy from the QoS configuration

Use the following syntax to delete a network policy.

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

7.2.3 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

The following is a sample configuration output.

```
A:ALA-12>config>qos# info detail
-----
...
    network 1 create
      description "Default network QoS policy."
      scope template
      ingress
      default-action fc be profile out
...
    network 600 create
      description "Default network QoS policy."
      scope template
      ingress
      default-action fc be profile out
...
    network 700 create
      description "Default network QoS policy."
      scope template
      ingress
      default-action fc be profile out
...
-----
A:ALA-12>config>qos#
```

7.2.4 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.

7.3 Network QoS policy command reference

- [Configuration commands for 7210 SAS-D and 7210 SAS-Dxp](#)
- [Operational commands](#)
- [Show commands](#)

7.3.1 Command hierarchies

7.3.1.1 Configuration commands for 7210 SAS-D and 7210 SAS-Dxp

```
- config
- qos
- [no] network network-policy-id [create]
- description description-string
- no description
- egress
- [no] fc fc-name
- [no] de-mark [force de-value]
- dot1p dot1p-priority
- no dot1p
- dot1p-in-profile dot1p-priority
- no dot1p-in-profile
- dot1p-out-profile dot1p-priority
- no dot1p-out-profile
- dscp-in-profile dscp-name
- no dscp-in-profile
- dscp-out-profile dscp-name
- no dscp-out-profile
- [no] remarking {use-dot1p | use-dscp | all}
- ingress
- default-action fc fc-name profile {in | out | use-dei}
- dot1p dot1p-priority fc fc-name profile {in | out}
- no dot1p dot1p-priority
- [no] fc fc-name [create]
- meter meter-id
- no meter
- multicast-meter meter-id
- no multicast-meter
- meter meter-id [multipoint] [create]
- no meter meter-id
- adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
- no adaptation-rule
- cbs size [kbits | bytes | kbytes]
- no cbs
- mbs size [kbits | bytes | kbytes]
- no mbs
- mode {trtcm1 | trtcm2 | srtcm}
- no mode
- rate cir-rate-in-kbps [pir pir-rate-in-kbps]
- no rate
- scope {exclusive | template}
- no scope
```

7.3.1.2 Operational commands

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

7.3.1.3 Show commands

```
- show
  - qos
    - network policy-id [detail]
    - network [network-policy-id] association
    - network [network-policy-id] [detail]
```

7.4 Command descriptions

7.4.1 Configuration commands

7.4.1.1 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.

7.4.1.2 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*. It 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. Specifies the source policy that the copy command will copy and the destination policy to which the command will duplicate the policy to a new or different policy ID.

Values 1 to 65535

overwrite

Specifies that 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.

remarking

Syntax

[no] **remarking** {**use-dot1p** | **use-dscp** | **all**}

Context

config>qos>network>egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command remarks egress packets sent out of access ports and access-uplink ports. Remarking can be enabled or disabled. On access port and access-uplink port egress, the behavior is as follows.

If remarking is enabled without specifying one of the options, by default 'use-dot1p' is used for access-egress and "all" is used for network-egress (that is, access-uplink port egress).

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

Default

no remarking

Parameters

use-dot1p

Specifies that the dot1p bits are marked in the packet header for all IEEE 802.1q and IEEE 802.1p encapsulated traffic sent out of the access port.

use-dscp

Specifies that the IP DSCP bits are marked in the packet header for IPv4 traffic sent out of the access port.

all

Specifies that the dot1p bits are marked in the packet header for all IEEE 802.1q and IEEE 802.1p encapsulated traffic, and in addition the IP DSCP bits are marked in the packet header for all IPv4 traffic sent out the access port.

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 to the default value.

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

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

7.4.1.3 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 defines the treatment packets receive as they ingress and egress the access uplink port and network IP interface in network mode of operation.

The QoS network policy consists of an ingress and egress component. The ingress component of the policy defines how dot1p bits are mapped to the internal FC and profile state. The FC and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the 7210 SAS. The 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 FC. Traffic mapped to each FC can be rate limited using separate meters for unicast and multipoint traffic.

The egress component of the network QoS policy defines the FC and profile to packet header priority bits (for example: dot1p bits).

The network policy 1 cannot be modified or deleted. It defines the default dot1p to FC mapping and default meters for unicast traffic and optional multipoint meters for BUM traffic on the, ingress. For the egress, it defines eight FCs which represent individual queues and the packet marking behavior.

If a new network policy is created (for instance, policy 2), only the default action and egress FC parameters are identical to the default policy. A new network policy does not contain the default dot1p to FC mapping for network QoS policy. The default network policy can be copied (use 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.

Changes made to an existing policy, using any of the sub-commands, are applied immediately to all access uplink ports where this policy is applied. For this reason, when many changes are required on a policy, Nokia highly recommends that the policy be copied to a work area. 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 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

create

Specifies that a QoS network policy is created.

7.4.1.4 Network ingress QoS policy commands

fc

Syntax

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

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 FC. After the *fc-name* is created, classification actions can be applied, which can be used in match classification criteria. By default, undefined FCs use the parameters configured in the default **policy** *policy-id* 1.

The **no** form of this command removes all the explicit meter mappings for *fc-name* forwarding types. The meter mappings revert to the default meters for *fc-name*.

Parameters

fc-name

Specifies the case-sensitive, system-defined FC name for which policy entries will be created.

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

create

Specifies that the FC is created. The **create** keyword requirement can be enabled in the **environment>create** context.

ingress

Syntax

ingress

Context

config>qos>network

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context create or edit policy entries that specify the dot1p to FC mapping for all packets.

When pre-marked packets ingress on a network port, the QoS treatment through the 7210 SAS is based on the mapping defined under the current node.

default-action

Syntax

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

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the default action taken for packets that have undefined dot1p bits set. The **default-action** command specifies the FC to which such packets are assigned.

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

Default

default-action fc be profile out

Parameters

fc *fc-name*

Specifies the FC name. All packets with dot1p or dot1p bits that are not defined are placed in this FC.

Values be | l2 | af | l1 | h2 | e | h1 | nc

profile {in | out}

Specifies an in or out of profile for all packets assigned to this FC. 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'.

fc

Syntax

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

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 FC. After the *fc-name* is created, classification actions can be applied and it can be used in match classification criteria. By default, undefined FCs use the parameters configured in the default network **policy** *policy-id* 1.

The **no** form of this command removes all the explicit meter mappings for *fc-name* forwarding types. The meter mappings revert to the default meters for *fc-name*.

Parameters

fc-name

Specifies the case-sensitive system-defined FC name for which policy entries are created.

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

create

Specifies that the FC is created. The **create** keyword requirement can be enabled in the **environment>create** context.

dot1p

Syntax

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

no dot1p *dot1p-priority*

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command explicitly sets the FC or enqueueing priority and profile of the packet when a packet has the *dot1p-priority* specified. Adding a dot1p rule on the policy forces packets that match the *dot1p-priority* specified to be assigned to the FC and profile of the packet based on the parameters included in the dot1p rule.

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

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

Parameters

dot1p-priority

Specifies the unique IEEE 802.1P value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-priority* value, the previous FC is overridden by the new parameters.

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

Values 0 to 7

fc fc-name

Specifies a value that must be one of the predefined FCs in the system. Specifying the *fc-name* is optional. When a packet matches the rule, the FC is only overridden when the **fc fc-name** parameter is defined on the rule. If the packet matches and the FC is not explicitly defined in the rule, the FC is inherited based on previous rule matches.

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

profile {in|out}

Specifies an in or out of profile for all packets assigned to this FC. 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'.

meter

Syntax

meter *meter-id* [**multipoint**] [**create**]

no meter *meter-id*

Context

config>qos>network>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an ingress network QoS policy meter. The **meter** command allows the creation of multipoint meters. Only multipoint meters can receive ingress packets that need to be sent to multiple destinations.

Multipoint meters are for traffic bound to multiple destinations. Within non-multipoint services, such as Epipe services, all traffic is considered unicast because of the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service will not be mapped to a multipoint service meter.

The **no** form of this command removes the meter-id from the network ingress QoS policy and from any existing ports using the policy. If any FC forwarding types are mapped to the meter, they revert to their default meters. When a meter is removed, any pending accounting information is discarded for each port meter created because of the definition of the meter in the policy is discarded.

Default

meter 1 (for unicast traffic)

meter 9 multipoint (for traffic other than unicast traffic)

Parameters

meter-id

Specifies the meter-id that uniquely identifies the meter within the policy. This is a required parameter each time the meter command is executed.

Values 1 to 12

multipoint

Specifies that this *meter-id* is for multipoint forwarded traffic only. This *meter-id* can only be explicitly mapped to the FC multicast, broadcast, or unknown unicast ingress traffic. If you attempt to map FC unicast traffic to a multipoint queue, an error is generated and no changes are made to the current unicast traffic queue mapping.

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

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

create

Keyword to create the ingress network QoS policy meter.

meter

Syntax

meter *meter-id*

no meter

Context

config>qos>network>ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default unicast forwarding type meter mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy as a non-multipoint meter before the mapping can be made. After the FC mapping is executed, all unicast traffic on a port using this policy is forwarded using the *meter-id*.

The **no** form of this command reverts the unicast (point-to-point) meter ID back to the default meter for the FC.

Default

meter 1

Parameters

meter-id

Specifies the meter ID. The specified parameter must be an existing, non-multipoint meter defined in the **config>qos>network>ingress** context.

Values 1 to 12

multicast-meter

Syntax

multicast-meter *meter-id*

no multicast-meter

Context

config>qos>network>ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default multicast forwarding type meter mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy as a multipoint meter before the mapping can be made. After the FC mapping is executed, all multicast traffic on a port using this policy is forwarded using the meter ID.

The **no** form of this command reverts the multicast forwarding type *meter-id* to the default meter for the FC.

Default

9

Parameters

meter-id

Specifies the multicast meter. The specified parameter must be an existing multipoint meter defined in the **config>qos>network>ingress** context.

Values 2 to 12

adaptation-rule

Syntax

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

no adaptation-rule

Context

config>qos>network>ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the method used by the system to derive the operational CIR and PIR rates when the meter is provisioned in hardware. For the **cir** and **pir** parameters, the system attempts to find the best operational rate depending on the defined constraint.

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

Default

adaptation-rule cir closest pir closest

Parameters

cir adaptation-rule

Specifies the adaptation rule and defines the constraints enforced to adapt the CIR rate defined using the **meter meter-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used to derive the operational CIR rate for the meter. When the **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 CIR value is equal to or less than the specified rate, taking into account the hardware step size. When **max** is defined, the operational CIR is the next multiple of 8 kbps (for 7210 SAS-D) equal to or less than the specified rate. For 7210 SAS-Dxp, see [Adaptation rule for meters on 7210 SAS-Dxp devices](#) for information about calculating the next multiple equal to or less than the specified rate.

min — Specifies that the operational CIR value is equal to or greater than the specified rate, taking into account the hardware step size. When **min** is defined, the operational CIR is the next multiple of 8 kbps (for 7210 SAS-D) equal to or greater than the specified rate. For 7210 SAS-Dxp, see [Adaptation rule for meters on 7210 SAS-Dxp devices](#) for information about calculating the next multiple equal to or greater than the specified rate.

closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. When **closest** is defined, the operational CIR is the next multiple of 8 kbps (for 7210 SAS-D) closest to the specified rate. For 7210 SAS-Dxp, see [Adaptation rule for meters on 7210 SAS-Dxp devices](#) for information about calculating the next multiple closest to the specified rate.

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 **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. When **max** is defined, the operational PIR is the next multiple of 8 kbps (for 7210 SAS-D) equal to or less than the specified rate. For 7210 SAS-Dxp, see [Adaptation rule for meters on 7210 SAS-Dxp devices](#) for

information about calculating the next multiple equal to or less than the specified rate.

min — Specifies that the operational PIR value is equal to or greater than the specified rate, taking into account the hardware step size. When **min** is defined, the operational PIR is the next multiple of 8 kbps (for 7210 SAS-D) equal to or greater than the specified rate. For 7210 SAS-Dxp, see [Adaptation rule for meters on 7210 SAS-Dxp devices](#) for information about calculating the next multiple equal to or greater than the specified rate.

closest — Specifies that the operational PIR value is equal to the closest specified rate, taking into account the hardware step size. When **closest** is defined, the operational PIR is the next multiple of 8 kbps (for 7210 SAS-D) closest to the specified rate. For 7210 SAS-Dxp, see [Adaptation rule for meters on 7210 SAS-Dxp devices](#) for information about calculating the next multiple closest to the specified rate.

cbs

Syntax

cbs size [kbits | bytes | kbytes]

no cbs

Context

config>qos>network>ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

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

Default

32 kbits

Parameters

size

Specifies the size parameter is an integer expression of the number of kilobits, kilobytes, or bytes reserved for the meter; for example, if a value of 100 kb is needed, then enter the value 100. The bucket size is rounded to the next highest 4096 bytes boundary.

Values kbits — 4 to 16384, default (7210 SAS-D)
 4 to 2146959, default (7210 SAS-Dxp)
 bytes — 512 to 2097152, default (7210 SAS-D)
 512 to 274810752, default (7210 SAS-Dxp)
 kbytes — 1 to 2048, default (7210 SAS-D)
 1 to 268369, default (7210 SAS-Dxp)

kbits

Specifies that the value is in kilobits.

bytes

Specifies that the value is in kilobytes.

kbytes

Specifies that the value is in bytes.

mbs

Syntax

mbs size [kbits|bytes|kbytes]

no mbs

Context

config>qos>network>ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default MBS for the meter. The maximum burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the MBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter-configured parameters.



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

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

Default

128 kbits

Parameters

size

Specifies the size parameter is an integer expression of the number of kilobits reserved for the meter; for example, if a value of 100 kb is needed, then enter the value 100. The bucket size is rounded off to the next highest 4096 bytes boundary.

Values kbits — 4 to 16384, default (7210 SAS-D)
 4 to 2146959, default (7210 SAS-Dxp)
 bytes — 512 to 2097152, default (7210 SAS-D)
 512 to 274810752, default (7210 SAS-Dxp)
 kbytes — 1 to 2048, default (7210 SAS-D)
 1 to 268369, default (7210 SAS-Dxp)

kbits

Specifies that the value is in kilobits.

bytes

Specifies that the value is in kilobytes.

kbytes

Specifies that the value is in bytes.

mode

Syntax

mode {trtcm1|srtcm}

no mode

Context

config>qos>network>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 (trTCM1) or Single Rate Three Color Marker (srTCM). The mode command can be executed at any time.

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

Default

trtcm1

Parameters

trtcm1

Specifies the policing algorithm defined in RFC2698 and meters the packet stream and marks its packets green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked yellow or green depending on whether it exceeds or does not exceed the CIR. The trTCM1 is useful, for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate. Two token buckets are used, the CBS bucket and the MBS bucket. Tokens are added to the buckets based on the CIR and PIR rates. The algorithm deducts tokens from both the CBS and the MBS buckets to determine a profile for the packet.

srtcm

Specifies that the mode is configured as a srTCM and meters a packet stream and marks the packets green, yellow, or red. Marking is based on a CIR and two associated burst sizes, a CBS and an MBS. A packet is marked green if it does not exceed the CBS, yellow if it does exceed the CBS but not the cir and red otherwise. The srTCM is useful, for example, for ingress policing of a service, where only the length, not the peak rate, of the burst determines service eligibility.

rate

Syntax

rate *cir* *cir-rate-in-kbps* [**pir** *pir-rate-in-kbps*]

no rate

Context

config>qos>network>ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the administrative PIR and CIR parameters for the meter.

The rate command can be executed at any time, altering the PIR and CIR rates for all meters created through the association of the network QoS policy with the *meter-id*. The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the pir-rate value.



Note:

The value of rates are represented in 1000 kilobits per second and bursts are represented as 1024 kilobits per second.

The **no** form of this command reverts all meter instances created with this *meter-id* to the default.

Default

rate 0 pir max

Parameters

cir cir-rate-in-kbps

Specifies that the default administrative CIR used by the meter will be overridden. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and values must be a positive integer.

The actual CIR rate is dependent on the meter's **adaptation-rule** parameters and the hardware.

Values 0 to 20000000, max (7210 SAS-Dxp)
0 to 4000000, max (7210 SAS-D)

pir pir-rate-in-kbps

Specifies the administrative PIR rate, in kilobits, for the meter. When this command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of max is assumed. When the **rate** command is executed, a PIR setting is optional.

Fractional values are not allowed and must be specified as a positive integer.

The actual PIR rate is dependent on the meter's adaptation-rule parameters and the hardware.

Values 0 to 20000000, max (7210 SAS-Dxp)
0 to 4000000, max (7210 SAS-D)

7.4.1.5 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

This command creates or edits egress policy entries that specify the FC to marking values map to be instantiated when this policy is applied to the access-uplink port.

The FC and profile state mapping to marking values (for example: IEEE 802.1p bits, etc.) bits mapping for all packets are 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 FC name. The FC name represents an egress queue. The **fc** *fc-name* represents a CLI parent node that contains subcommands or parameters describing the marking criteria of packets flowing through it. The **fc** command overrides the default parameters for that FC to the values defined in the network default policy.

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

Default

undefined FCs default to the configured parameters in the default network policy (policy ID 1)

Parameters

fc-name

Specifies the case-sensitive, system-defined FC name for which policy entries will be created.

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

adaptation-rule

Syntax

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

no adaptation-rule

Context

config>qos>network>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. See [Table 17: Supported hardware rates and CIR and PIR values for egress queues on the 7210 SAS-D](#) and [Table 18: Supported hardware rates and CIR and PIR values for egress queues on the 7210 SAS-Dxp](#) for information about supported hardware step-size rates for 7210 SAS-D and 7210 SAS-Dxp.

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. See [Table 17: Supported hardware rates and CIR and PIR values for egress queues on the 7210 SAS-D](#) and [Table 18: Supported hardware rates and CIR and PIR values for egress queues on the 7210 SAS-Dxp](#) for information about supported hardware step-size rates for 7210 SAS-D and 7210 SAS-Dxp.

- 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.

slope-policy

Syntax

[no] **slope-policy** *name*

Context

config>qos>network>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 will be used by the queue. The non-ring slopes are used by the traffic received on access SAP ingress and sent out of access SAP egress queues. The ring slopes are used by the traffic received on access-uplink port ingress and sent out of access-uplink port egress queues.

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 the command reverts the default slope policy to the queue.

Parameters

name

Specifies an existing slope policy name. If the slope policy name does not exist, the **slope-policy** command fails.

Values 32 chars maximum

queue

Syntax

queue *queue-id* **create**

no queue

Context

config>qos>network

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates the context to modify queue parameters associated with a particular queue.

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

Parameters

queue-id

Specifies the ID of the queue.

Values 1 to 8

create

Specifies that a network queue policy will be created.

rate

Syntax

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

no rate

Context

config>qos>network>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 `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 the command reverts all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR 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 of 0 is assumed. Fractional values are not allowed and must be specified as a positive integer.

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

Values 0 to 100

Default 0

pir *pir-percent*

Specifies the percentage of the maximum rate allowed for the queue. When the **rate** command is executed, the PIR setting is optional. When the **rate** command has not been executed, or the PIR parameter is not explicitly specified, the default PIR of 100 is assumed. Fractional values are not allowed and must be specified as a positive integer.

Values 1 to 100 %

Default 100

7.4.1.6 Network egress QoS policy forwarding class commands

dot1p-in-profile

Syntax

dot1p-in-profile *dot1p-priority*

no dot1p-in-profile

Context

config>qos>network>egress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

The command enables the context to mark on an egress the in and out of profile status through a specific dot1p combination (similar to DEI option). It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with in-profile status (or green color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, this command has no effect.

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

Parameters

dot1p-priority

Specifies the unique IEEE 802.1P value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous FC and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a FC or enqueueing priority parameter is missing. A maximum of eight dot1p rules are allowed on a single policy.

Values 0 to 7

dot1p-out-profile

Syntax

dot1p-out-profile *dot1p-priority*

no dot1p-out-profile

Context

config>qos>network>egress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

The command enables the context to mark on an egress the in and out of profile status via a specific dot1p combination (similar to DEI option). It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with out-profile status (or yellow color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, this command has no effect.

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

Parameters

dot1p-priority

Specifies the unique IEEE 802.1P value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous FC and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a FC or enqueueing priority parameter is missing. A maximum of eight dot1p rules are allowed on a single policy.

Values 0 to 7

dscp-in-profile

Syntax

dscp-in-profile *dscp-name*

no dscp-in-profile

Context

config>qos>network>egress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the in-profile DSCP name for the FC. When marking is set, the corresponding DSCP value is used to mark all IP packets with in-profile status, on the egress of this FC queue.

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

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

Parameters

dscp-name

Specifies a system-defined or a user-defined, case-sensitive *dscp-name*.

dscp-out-profile

Syntax

dscp-out-profile *dscp-name*

no dscp-out-profile

Context

config>qos>network>egress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the out-of-profile DSCP name for the FC. 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 FC queue.

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

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

Parameters

dscp-name

Specifies a system-defined or a user-defined, case-sensitive *dscp-name*.

7.4.1.7 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 network policy information.

Parameters

policy-id

Displays information for the specific policy ID.

Values 1 to 65535

Default all network policies

detail

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

Output

The following outputs are examples of network policy information, and [Table 32: Output fields: QoS network](#) displays the output fields:

- [Sample output for 7210 SAS-D](#)
- [Sample output for 7210 SAS-D \(detail\)](#)

Sample output for 7210 SAS-D

```
*A:SN12345678# show qos network 1
=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id      : 1                      Remark      : True
Forward Class  : be                    Profile      : Out
Attach Mode    : l2                    Config Mode  : l2
Scope          : Template
Description    : Default network QoS policy.
-----
Meter Mode    CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS      MBS
-----
1      TrTcm_CA  0          closest    max        closest  32 KBytes 128 KBytes
9      TrTcm_CA  0          closest    max        closest  32 KBytes 128 KBytes
-----
FC              UCastM      MCastM
-----
No FC-Map Entries Found.
-----
Port Attachments
-----
Port-id : 1/1/24
=====
*A:SN12345678#

*A:dut-g# show qos network 1 detail
=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id      : 1                      Remark      : True
Forward Class  : be                    Profile      : Out
Attach Mode    : l2                    Config Mode  : l2
Scope          : Template
Description    : Default network QoS policy.
-----
Meter Mode    CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS      MBS
-----
1      TrTcm_CA  0          closest    max        closest  32        128
9      TrTcm_CA  0          closest    max        closest  32        128
-----
FC              UCastM      MCastM
-----
No FC-Map Entries Found.
-----
```

```

Dot1p Bit Map                Forwarding Class                Profile
-----
0                             be                             Out
1                             l2                             In
2                             af                             Out
3                             af                             In
4                             h2                             In
5                             ef                             In
6                             h1                             In
7                             nc                             In
-----

Egress Forwarding Class Queuing
-----
FC Value      : 0                FC Name      : be
- Dot1p Mapping
Out-of-Profile : 0                In-Profile   : 0

FC Value      : 1                FC Name      : l2
- Dot1p Mapping
Out-of-Profile : 1                In-Profile   : 1

FC Value      : 2                FC Name      : af
- Dot1p Mapping
Out-of-Profile : 2                In-Profile   : 3

FC Value      : 3                FC Name      : l1
- Dot1p Mapping
Out-of-Profile : 2                In-Profile   : 3

FC Value      : 4                FC Name      : h2
- Dot1p Mapping
Out-of-Profile : 4                In-Profile   : 4

FC Value      : 5                FC Name      : ef
- Dot1p Mapping
Out-of-Profile : 5                In-Profile   : 5

FC Value      : 6                FC Name      : h1
- Dot1p Mapping
Out-of-Profile : 6                In-Profile   : 6

FC Value      : 7                FC Name      : nc
- Dot1p Mapping
Out-of-Profile : 7                In-Profile   : 7
-----

Port Attachments
-----
Port-id : 1/1/24
=====
*A:dut-g#

```

Sample output for 7210 SAS-D (detail)

```

*A:K-SASD>config>qos# show qos network
Policy-id  Remark      LerUseDscp  Description
-----
1          False      False      Default network QoS policy.
19         True       False      Description for Network Policy id # 19
100        True       False
200        True       False
300        True       False

```

```

-----
*A:K-SASD>config>qos# show qos network 1 detail
=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id       : 1
Egr Remark      : False
Forward Class   : be                      Profile      : Out
Scope          : Template                 Policy Type   : port
Accounting      : packet-based
Description     : Default network QoS policy.

-----
DSCP (Egress)           Forwarding Class           Profile
-----
No Matching Entries

-----
Prec (Egress)           Forwarding Class           Profile
-----
No Matching Entries

-----
Dot1p Bit Map           Forwarding Class           Profile
-----
0                        be                      Out
1                        l2                      In
2                        af                      Out
3                        af                      In
4                        h2                      In
5                        ef                      In
6                        h1                      In
7                        nc                      In

-----
Meter Mode    CIR Admin CIR Rule  PIR Admin  PIR Rule  CBS Admin MBS Admin
              CIR Oper              PIR Oper              CBS Oper  MBS Oper
-----
1    TrTcm1_CA  0          closest    max        closest   def        def
              0
9    TrTcm1_CA  0          closest    max        closest   def        def
              0

-----
FC          UCastM      MCastM
-----
No FC-Map Entries Found.

-----
Egress Forwarding Class Queuing
-----
FC Value      : 0                      FC Name      : be
- DSCP Mapping
Out-of-Profile : be                      In-Profile   : be
- Dot1p Mapping
Out-of-Profile : 0                      In-Profile   : 0

```

| | |
|-----------------------|-------------------|
| FC Value : 1 | FC Name : l2 |
| - DSCP Mapping | |
| Out-of-Profile : cs1 | In-Profile : cs1 |
| - Dot1p Mapping | |
| Out-of-Profile : 1 | In-Profile : 1 |
| FC Value : 2 | FC Name : af |
| - DSCP Mapping | |
| Out-of-Profile : af12 | In-Profile : af11 |
| - Dot1p Mapping | |
| Out-of-Profile : 2 | In-Profile : 3 |
| FC Value : 3 | FC Name : l1 |
| - DSCP Mapping | |
| Out-of-Profile : af22 | In-Profile : af21 |
| - Dot1p Mapping | |
| Out-of-Profile : 2 | In-Profile : 3 |
| FC Value : 4 | FC Name : h2 |
| - DSCP Mapping | |
| Out-of-Profile : af41 | In-Profile : af41 |
| - Dot1p Mapping | |
| Out-of-Profile : 4 | In-Profile : 4 |
| FC Value : 5 | FC Name : ef |
| - DSCP Mapping | |
| Out-of-Profile : ef | In-Profile : ef |
| - Dot1p Mapping | |
| Out-of-Profile : 5 | In-Profile : 5 |
| FC Value : 6 | FC Name : h1 |
| - DSCP Mapping | |
| Out-of-Profile : nc1 | In-Profile : nc1 |
| - Dot1p Mapping | |
| Out-of-Profile : 6 | In-Profile : 6 |
| FC Value : 7 | FC Name : nc |
| - DSCP Mapping | |
| Out-of-Profile : nc2 | In-Profile : nc2 |
| - Dot1p Mapping | |
| Out-of-Profile : 7 | In-Profile : 7 |

Port Attachments

Port-id : 1/1/10

=====

Table 32: Output fields: QoS network

| Label | Description |
|---------------------------|---|
| Policy-Id | Displays the ID that uniquely identifies the policy. |
| Remark | True — Remarking is enabled. False — Remarking is disabled. |
| Description | Displays a text string that helps identify the policy's context in the configuration file. |
| Forward Class/ FC Name | Specifies the FC name. |
| Profile | Out — Specifies the EXP marking for the packets that are out-of-profile, egressing on this queue.Specifies the dot1p marking for the packets which are out-of-profile, egressing on this queue |
| | In — Specifies the EXP marking for the packets that are in-of-profile, egressing on this queue.Specifies the dot1p markings for in-profile packets egressing this queue. |
| Accounting | Packet-based — Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow Frame-based — Specifies that the meters associated with this policy account for the packet framing overheads (such as for Ethernet the IFG and preamble), while accounting the bandwidth to be used by the flow |
| dot1p Bit Mapping: | |
| Out-of-Profile | Displays the dot1p value used for out-of-profile traffic |
| In-Profile | Displays the dot1p value used for in-profile traffic |
| Port Attachment | |
| Port-Id | Specifies the physical port identifier that associates the interface |

8 Network queue QoS policies

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

8.1 Overview

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.

On 7210 SAS-D and 7210 SAS-Dxp, all network-queue policies contain eight queues. The user cannot allocate fewer queues or change the mapping of FC to queue. Each of these queues is 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:

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.

On 7210 SAS-D and 7210 SAS-Dxp, FCs are mapped to eight queues available at the port according to [Table 15: FC-to-queue ID map](#).

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]
```

```
*A:7210SASD>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
-----
```

*7210SASD>config>qos>network-queue#

*7210SASDxp>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 10 pir 100
    adaptation-rule cir closest pir closest
exit
queue 3
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
exit
queue 4
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
exit
queue 5
    rate cir 15 pir 100
    adaptation-rule cir closest pir closest
exit
queue 6
    rate cir 15 pir 100
    adaptation-rule cir closest pir closest
-----
```

```
exit
queue 7
    rate cir 5 pir 100
    adaptation-rule cir closest pir closest
exit
queue 8
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
exit
```

8.2.2 Applying network queue configuration in access-uplink mode

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

```
config>port#
    ethernet
        access
            uplink
                queue-policy policy-name

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

8.3 Default network queue policy values

The default network queue policies are identified as *policy-id* **default**. The default policies cannot be modified or deleted. The following table lists the network queue policy defaults for 7210 SAS-D.

Table 33: Network queue policy defaults for 7210 SAS-D

| Field | Default |
|-------------|-----------------------------------|
| description | Default network queue QoS policy. |
| queue 1 | |
| rate | |
| cir | 0 |
| pir | 100 |

| Field | Default |
|---------|---------|
| queue 2 | |
| rate | |
| cir | 25 |
| pir | 100 |
| queue 3 | |
| rate | |
| cir | 25 |
| pir | 100 |
| queue 4 | |
| rate | |
| cir | 25 |
| pir | 100 |
| queue 5 | |
| rate | |
| cir | 100 |
| pir | 100 |
| queue 6 | |
| rate | |
| cir | 100 |
| pir | 100 |
| queue 7 | |
| rate | |
| cir | 10 |
| pir | 100 |
| queue 8 | |
| rate | |
| cir | 10 |

| Field | Default |
|-------|---------|
| pir | 100 |

The following table lists the network queue policy defaults for 7210 SAS-Dxp.

Table 34: Network queue policy defaults for 7210 SAS-Dxp

| Field | Default |
|-------------|-----------------------------------|
| description | Default network queue QoS policy. |
| queue 1 | |
| rate | |
| cir | 0 |
| pir | 100 |
| queue 2 | |
| rate | |
| cir | 10 |
| pir | 100 |
| queue 3 | |
| rate | |
| cir | 10 |
| pir | 100 |
| queue 4 | |
| rate | |
| cir | 10 |
| pir | 100 |
| queue 5 | |
| rate | |
| cir | 15 |
| pir | 100 |
| queue 6 | |
| rate | |

| Field | Default |
|---------|---------|
| cir | 15 |
| pir | 100 |
| queue 7 | |
| rate | |
| cir | 5 |
| pir | 100 |
| queue 8 | |
| rate | |
| cir | 10 |
| pir | 100 |

The following displays network queue default policy parameters for 7210 SAS-D.

```
*A:dut-a>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
```

The following displays network queue default policy parameters for 7210 SAS-Dxp.

```
*A:dut-a>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 10 pir 100
    adaptation-rule cir closest pir closest
exit
queue 3
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
exit
queue 4
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
exit
queue 5
    rate cir 15 pir 100
    adaptation-rule cir closest pir closest
exit
queue 6
    rate cir 15 pir 100
    adaptation-rule cir closest pir closest
exit
queue 7
    rate cir 5 pir 100
    adaptation-rule cir closest pir closest
exit
queue 8
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
exit
```

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**.

Use the following syntax 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

The following example shows the command usage to copy and overwrite a QoS policy.

```
config>qos# copy network-queue nq1-cbs nq2-cbs
```

Example

The following is a sample 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

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
- queue queue-id
- adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
- no adaptation-rule
- rate cir-rate-in-kbps [pir pir-rate-in-kbps]
- no rate
```

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]
```

8.6 Command descriptions

8.6.1 Configuration commands

8.6.1.1 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.

8.6.1.2 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. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters

network-queue *src-name dst-name*

Specifies that the source policy ID and the destination policy ID are network-queue policy IDs. Specifies the source policy ID that the **copy** command attempts to copy from, and the destination policy ID to which the command copies a duplicate of the policy.

overwrite

Specifies that everything in the existing destination policy is overwritten with the contents of the source policy. If **overwrite** is not specified, a message is generated saying that the destination policy ID exists.

8.6.1.3 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 creates a context to configure a network queue policy. Network queue policies on the Ethernet port define network egress queuing.

Network queue policies define the egress queuing for access-uplink ports.

The **no** form of this command removes the network-queue policy from use. However, the network queue with *policy-name* **default** cannot be modified or deleted.

Default

default

Parameters

policy-name

Specifies the name of the network queue policy. Allowed values are 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

Specifies that a network queue policy is created.

8.6.1.4 Network queue QoS policy queue commands

queue

Syntax

queue *queue-id*

Context

config>qos>network-queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a QoS network-queue policy queue.

The FCs are mapped to queues, as listed in [Table 15: FC-to-queue ID map](#). Only one FC can be mapped to one queue. A *queue-id* value of 8 is the highest priority and a *queue-id* value of 1 is the lowest priority.

Network queues carry both the unicast and multicast traffic and no segregation is performed. The hardware port scheduler prioritizes the queue according to the priority for each queue. High priority traffic should be mapped to high-priority FCs. Mapping traffic to high-priority FCs does not guarantee high priority treatment, because the scheduler policy can influence the relative priority among the queues.

Parameters

queue-id

Specifies the queue within the policy. This is a required parameter each time the queue command is executed.

Values 1 to 8

adaptation-rule

Syntax

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

no adaptation-rule

Context

config>qos>network-queue>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

The **no** form of 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 **pir** 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.

slope-policy

Syntax

[no] **slope-policy** *name*

Context

config>qos>network-queue>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 will fail. If a slope policy is currently associated with a queue, the slope policy cannot be removed from the system.

The slope policy contains High and Low slope definitions that will be used by the queue.

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 the command restores the default slope policy to the queue.

Parameters

name

The name parameter is required and must specify an existing slope policy name. If slope-policy-name does not exist, the slope-policy command will fail.

Values 32 chars maximum

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 PIR and the CIR parameters for the queue. The PIR defines the maximum rate at which 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 any time, changing 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 and CIR 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 of 0 is assumed. Fractional values are not allowed and must be specified as a positive integer.

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

Values 0 to 100

Default 0

pir percent

Specifies the percentage of the maximum rate allowed for the queue. When the **rate** command is executed, the PIR setting is optional. When the **rate** command has not been executed, or the PIR parameter is not explicitly specified, the default PIR of 100 is assumed. Fractional values are not allowed and the value must be a positive integer.

Values 1 to 100 %

Default 100

8.6.1.5 Show commands

network-queue

Syntax

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

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays network queue policy information.

Parameters

network-queue-policy-name

Specifies the name of the network queue policy. Allowed values are 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.

detail

Specifies 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 policy information, and [Table 35: Output fields: network queue policy](#) describes the output fields.

Sample output

```
*A:card-1# show qos network-queue nql
=====
QoS Network Queue Policy
-----
Network Queue Policy (nql)
-----
Policy          : nql
Accounting      : packet-based
-----
Associations
-----
Port-id : 1/1/20
=====
*A:card-1#

*A:card-1>config>qos# show qos network-queue nql-cbs detail
=====
QoS Network Queue Policy
-----
Network Queue Policy (nql-cbs)
-----
Policy : nql-cbs
Accounting : packet-based
-----
Queue CIR      PIR      CBS
      CIR Rule  PIR Rule
-----
1      0      32      8.29
      max      closest
2      0      100     6.00
      closest  closest
3      0      100    10.00
      closest  closest
4      0      100     6.00
      closest  closest
5      0      100    10.00
      closest  closest
6      0       4    10.00
      closest  closest
7      3      93     1.00
      closest  closest
8      0       3     7.00
      closest  closest
-----
FC      UCastQ
-----
be      1
l2      2
af      3
l1      4
h2      5
ef      6
h1      7
nc      8
-----
Associations
-----
Port-id : 1/1/1
=====
*A:card-1>config>qos#
```

Table 35: Output fields: network queue policy

| 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 has two classification resources and one meter (the **num-qos-classifiers** set to value "2"). Service ingress queuing is not supported on the 7210 SAS-D and 7210 SAS-Dxp. Instead the user has the option of using policing per FC (and for VPLS services, per FC and per traffic type). The default policies can be copied and modified but cannot be deleted. The default policies are identified as policy ID 1.

The default policies are applied to the appropriate interface, by default. For example, the default SAP ingress policy is applied to access ingress SAPs. Other QoS policies must be explicitly associated. See "CLI Usage" in the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide for more information about the tasks and commands necessary to access the CLI, and to configure and maintain devices.

9.1.1 Default SAP ingress policy

The default policy 1 maps all traffic to default FC "be" and maps FC "be" to meter 1. Meter 1 is configured with cir 0 and pir max.

Example

The following is a sample configuration output.

```
*A:7210-SAS>config>qos>sap-ingress# info detail
-----
description "Default SAP ingress QoS policy."
num-qos-classifiers 2
scope template
meter 1 create
  mode trtcml
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
  mbs default
  cbs default
exit
default-fc "be"
-----
*A:7210-SAS>config>qos>sap-ingress#
```

9.1.1.1 SAP ingress policy defaults

The following table lists the SAP ingress policy defaults.

Table 36: SAP ingress policy defaults

| Field | Default |
|---------------------|-----------------------------------|
| description | "Default SAP ingress QoS policy." |
| scope | template |
| num-qos-classifiers | 2 |
| meter | 1 |
| mode | trtcm1 |
| adaptation-rule | cir closest pir closest |
| rate | pir = max, cir= 0 |
| cbs | 32kbits |
| mbs | 128kbits |
| default-fc | be |

9.1.1.2 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 the reboot can be recreated during 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 attempts to allocate resources for all the SAPs on a first-come-first-served basis. There is no guarantee that resources will be allocated to all SAPs. When the index file is not present, it is possible that saved the configuration does not execute successfully after the reboot.

**Note:**

The index file used for QoS maps is different from the one used for storing interface indexes.

9.1.1.2.1 Use of the keyword multipoint for default meter "11"

The system allows the sharing of a single meter for both unicast and multipoint traffic. The user can configure any of the available meters for multipoint traffic. The use of the **multipoint** keyword during meter creation is deprecated, except for use with meter 11 as described in the following paragraphs.

When the **multipoint** keyword is specified with meter 11 the software interprets it to be the default multipoint meter. The default multipoint meter is used for all FCs that do not have explicit multipoint meters configured. The software performs resource checks to ensure that resources needed to use the multipoint meter with all the FCs are available before allowing this change.

**Note:**

- When the **num-qos-resources** command is set to a value of 2, the default multipoint meter 11 cannot be used because only a single meter is available for use.
- When associating a meter with an FC for BUM traffic, the software does not validate whether the meter is a multipoint meter, which allows the user to use a single meter for unicast and BUM traffic. This implies efficient use of SAP ingress qos resources. From Release 4.0R4 onwards when the **multipoint** keyword is used, the software issues a warning indicating that it is an obsolete CLI command, and it is not saved in the configuration file, deprecating the use of multipoint keyword with any meter other than the default.

9.1.1.2.1.1 Example uses of the multipoint meter

Example: 1

```
*7210-SAS>config>qos# sap-ingress 12 create
*7210-SAS>config>qos>sap-ingress$ info
-----
num-qos-classifiers 4
meter 1 create
exit
-----
*7210-SAS>config>qos>sap-ingress$
```

In example 1, all FCs in the SAP ingress policy use the default meter 1 (for all traffic types). If the **configure qos sap-ingress id meter 11 multipoint create** command is executed, it attaches the default meter "11" with all the FCs defined in the SAP ingress policy.

After this configuration, all the FCs in this policy use two meters: default meter "1" to meter unicast traffic for all the FCs, and meter "11" to meter BUM traffic for all the FCs. In this example, because only the default FC "be" is in use, the multipoint meter is used to meter BUM traffic associated with default FC "be".

The following example shows the policy after the configuration is changed.

```
*7210-SAS>config>qos# sap-ingress 12
*7210-SAS>config>qos>sap-ingress$ info
-----
num-qos-classifiers 4
meter 1 create
exit
meter 11 multipoint create
-----
*7210-SAS>config>qos>sap-ingress$
```

Delete the multipoint meter "11" to remove all the FCs associated with the multicast-meter (assuming all the FCs are using the default multicast meter and do not have any other multicast meter explicitly configured). Execute the **configure qos sap-ingress id no meter 11** command to dissociate meter "11" from the FCs, and the FCs use only meter "1" (if no other meter configured explicitly).

Example: 2

```
*7210-SAS>config>qos# sap-ingress 12
*7210-SAS>config>qos>sap-ingress$ info
-----
configure> qos> sap-ingress 10 create
```

```

meter 1 create
exit
meter 3 create
exit
default-fc be
fc be
meter 3
multicast-meter 3
exit
fc af
meter 3
exit
exit
-----

```

Starting with the policy in example 2, if the user executes the **configure qos sap-ingress <id> meter 11 multipoint create** command, the FC "be" continues to use meter "3" and the FC "af" uses meter "11" for BUM traffic. In this example, if the user executes the **configure qos sap-ingress id fc be no multicast-meter** command, the default meter "11" is used for FC "be".

Example: 3

```

-----
configure> qos sap-ingress 10 create
meter 1 create
exit
meter 3 create
exit

default-fc be

fc be
meter 3
unknown-meter 3
exit
exit
-----

```

On execution of the **configure qos sap-ingress id meter 11 multipoint create** command, the FC "be" unknown-unicast traffic type continues to use meter 3 and broadcast and multicast traffic types use meter "11".

In example 3, if a broadcast-meter was initially configured in the SAP-ingress policy and then followed by execution of the **configure qos sap-ingress id meter 11 multipoint create** command, the FC "be" changes to use meter "11" for multicast traffic, and broadcast traffic continues to use meter "3" for unknown-unicast traffic and meter "3" for unicast traffic.

In example 3, if the user executes **configure qos sap-ingress id fc be no unknown-meter** command, meter "3" is used for all traffic types classified to FC "be". But, if the default meter "11" is defined in the policy, FC "be" uses meter "11" for BUM traffic.

9.1.1.3 Service ingress meter selection rules

This section describes the rules for meter selection by different traffic types under various configurations for VPLS services:

In the default policy, only meter "1" is defined. All FC and all traffic types use meter "1" by default. Meter "11" is not created by default and is not available for use.

Example

The following is a sample default policy configuration output.

```
*7210-SAS>config>qos# sap-ingress 1 create // Default policy
*7210-SAS>config>qos>sap-ingress$ info
-----
num-qos-classifiers 2
meter 1 create
exit
-----
*7210-SAS>config>qos>sap-ingress$
```

The following describes the use of meters in a VPLS service, when meter "11" is not configured in the policy:

- If an FC is created without explicit meters, the default meter "1" is used for unicast traffic and for multipoint traffic types (such as broadcast, multicast, and unknown-unicast traffic).
- If an FC is created with an explicit unicast meter, that meter is used for unicast traffic and for multipoint traffic types (such as broadcast, multicast, and unknown-unicast traffic).
- If an FC is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively, and use the unicast meter for all other traffic types.
- If an FC is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other traffic types.
- If an FC is created with an explicit unicast meter, an explicit broadcast meter, and an explicit multicast meter, use these meters for unicast, broadcast, and multicast traffic types respectively. The unknown unicast traffic types use the explicitly defined multicast meter.
- If an FC is created with an explicit unicast meter, explicit broadcast meter, explicit unknown-unicast meter, and explicit multicast meter, use these meters for unicast, broadcast, unknown-unicast, and multicast traffic types respectively.

The following describes the use of meters in a VPLS service, when meter "11" is defined in the policy:

- If an FC is created without explicit meters, use the default meter "1" for unicast traffic and default meter "11" for all other traffic types (such as broadcast, multicast, and unknown-unicast).
- If an FC is created with an explicit unicast meter, use that meter for unicast traffic, and use default meter "11" for all other traffic types.
- If an FC is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively, and use meter "11" for all other traffic types.
- If an FC is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic.
- If an FC is created with an explicit unicast meter, explicit broadcast meter, and explicit multicast meter, use these meters for unicast, broadcast, and multicast traffic types respectively. Unknown unicast traffic types use the explicitly defined multicast meter.
- If an FC is created with an explicit unicast meter, explicit broadcast meter, explicit unknown-unicast meter, and explicit multicast meter, use these meters for unicast, broadcast, unknown-unicast, and multicast traffic types respectively.

The following are rules for meter selection for Epipe and IES services:



Note:

These rules apply to IES services when PIM, multicast is not enabled in the service.

- IPv4 multicast with PIM is not supported on all 7210 SAS platforms. Check the 7210 SAS release notes to know the availability of this feature on all platforms.
- A multipoint meter cannot be used. A multipoint meter configured in a policy is not used when the policy is applied to a SAP in an Epipe service.
- All FCs associated with a meter always use the unicast meter.

9.1.1.4 Service ingress QoS policy configuration considerations

The **num-qos-classifiers** command cannot be modified when the policy is in use (for example, when it is associated with a SAP). Other parameters in the SAP ingress policy can be changed.

When changing other parameters (for example, FC meter map or FC classification match criteria entries) for a policy that is in use, the system recomputes the resources required to accommodate the change. If the resources required exceed the configured value for **num-qos-classifiers**, the change is not allowed.

If more resources are needed than are configured in **num-qos-classifiers** for an existing policy, the following options are available:

- Copy the existing policy to a new policy, modify the **num-qos-classifiers** command, modify the match criteria entries, and finally modify the SAP configuration to associate it with the new policy.
- Ensure the existing policy is not in use by any SAP (if required, change the SAP configuration to disable the use of the QoS policy using the **no qos** form of the command), change all the required parameters, and finally modify the SAP configuration to use the policy again.



Note:

Both these options have side-effects; for example, they can reset the statistics associated with the meters and potentially cause existing traffic classification not to take effect; however, the system ensures that the default policy is in use during the intermittent time when the policy changes are being made after the two options are performed.

- In releases before Release 3.0R1, the software always computes the number of resources (like classifiers and meters) required by a policy assuming it will be used in a VPLS service. This allows the policy to be applied to either an Epipe or VPLS service.
- From Release 3.0R1 onwards, when creating a SAP ingress policy, the software does not compute the number of resources required by a policy and validate it against resources available in the system. The software validates the resources needed only when the SAP ingress policy is attached to a SAP. If enough resources are available, the association succeeds, otherwise the command fails. Based on the service (VLL, VPLS, and so on) in which the SAP is configured, for the same SAP ingress policy the amount of resources required is different. The software validates that the number of QoS resources specified with the **num-qos-classifiers** command is sufficient for the match criteria, FC, and service specified, and the resources are available in hardware. If the validation fails, the software disallows the association of the SAP ingress policy with the SAP.
- The match criteria type (that is, **mac-criteria**, **ipv4-criteria**, and **ipv6-criteria**) cannot be changed when the SAP ingress QoS policy is in use. For example, if the match criteria type is set to **ipv4-criteria** and the policy is associated with a SAP, the **ipv6-criteria** or **mac-criteria** cannot be enabled in the same policy. If there is a need to change the criteria, the user must remove the association and then change the SAP ingress policy to use the new match criteria.

See [Resource allocation for service ingress QoS policy classification rules](#) for more information.

9.1.2 Resource allocation for service ingress QoS policy classification rules

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 and meter resources for use from this pool. Further on, resources can be allocated for different SAP ingress QoS policy classification match criteria, based on the operator needs. Users can modify 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 fixed number of entries.

The number of slices allotted for a SAP ingress QoS policy is specified using the **configure system resource-profile ingress-internal-tcam qos-sap-ingress-resource** command.

User can specify a limit for the amount of resources required for SAP ingress QoS policies and also an 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, or IPv6 criteria. Allocation of classification entries also allocates meter/policer resources, used to implement per FC per traffic type policing.

By default, the system allocates resources for SAP ingress QoS policies to maintain backward compatibility with Release 4.0 and allocates resources for MAC criteria and IP criteria (by setting it to 'max'). Setting the value to 'max' allows each match criterion to use the available SAP ingress QoS resources on a first-come-first-served model. By default, the software does not allocate resources for use by ingress IPv6 filters. Before associating an IPv6 SAP ingress policy to a SAP, resources must be allocated. Until resources are allocated for use by IPv6 filters, software fails all attempts to associate an IPv6 filter policy with a SAP.

When the user allocates resources for use by SAP ingress QoS policies using the **configure system resource-profile qos-sap-ingress-resource** command, the system allocates resources in slices of 256 entries. The following table describes the use of these entries by different types of match criteria.

**Note:**

See the *7210 SAS Release Notes* for information about services supported on different 7210 SAS platforms. The references to services in the following table appear for completeness and do not imply support is available.

Table 37: SAP ingress resource allocation and match criteria types

| Type of match criteria | Description |
|------------------------|--|
| mac-criteria (any) | <p>User needs to allocate resources for mac-criteria from the SAP ingress QoS resource pool by using the configure system resource-profile ingress-internal-tcam qos-sap-ingress-resource mac-match-enable command before using SAP ingress policies with mac-criteria. Every entry configured in the SAP ingress QoS policy using the mac-criteria uses one entry from the slices in the hardware.</p> <p>For example, assume a SAP-ingress QoS policy is configured to use mac-criteria with 25 entries and uses configure system resource-profile ingress-internal-tcam qos-sap-ingress-resource mac-match-enable 1, to configure one slice for use by mac-criteria (allowing a total of 256 entries for use by policies using mac-criteria). In this case, the user can have 10 SAPs using mac-criteria SAP ingress policy and consumes 250 entries.</p> |
| ipv4-criteria (any) | <p>The usage is same as the mac-criteria. Resources need to be allocated using the configure system resource-profile ingress-internal-tcam qos-sap-ingress-resource ipv4-match-enable command. Additionally, IPv4 criteria can share the</p> |

| Type of match criteria | Description |
|--|---|
| | <p>entries allocated for IPv6 criteria. The software automatically allocates entries from an IPv6 criteria slice to IPv4 criteria policies, if there are no entries available in the allocated IPv4 criteria slices and there are no slices available for allocation to IPv4 criteria from the SAP ingress QoS resource pool. The number of hardware entries taken up by an IPv4 criteria entry when using the IPv6 criteria slices is the same as required by an entry using IPv6 criteria (see below for details).</p> |
| ipv6-criteria (any) | <p>User needs to allocate resources from the SAP ingress QoS resource pool for ipv6-criteria by using the configure system resource-profile ingress-internal-tcam qos-sap-ingress-resource ipv6-ipv4-match-enable command before using IPv6 criteria and num-qos-classifiers must specify the ipv6 keyword. Every ipv4 criteria match entry or ipv6 criteria match entry configured in the QoS policy using ipv6-criteria uses two entries from the slices allocated for use by ipv6-criteria (128-bit) in the hardware. Software allocates entries from the ipv6-criteria pool if the SAP ingress QoS policy uses both ipv6-criteria entries and ipv4-criteria (any or IPv4 DSCP) entries or if the SAP ingress QoS policy uses only IPv6 criteria any or if the SAP ingress QoS policy uses ipv4 criteria any and there are no resources available in the IPv4 criteria (as described above).</p> <p>For example, assume a QoS policy is configured to use ipv6-criteria with 25 entries and using configure system resource-profile ingress-internal-tcam qos-sap-ingress-resource ipv4-ipv6-128-match-enable 1, user configures one slice for use by ipv6-criteria. This allows for a total of 128 entries for use by SAPs using SAP ingress QoS policies with ipv6-criteria (as each IPv6 entry uses two entries in hardware). In this example, user can have five SAPs using this policy and consuming 125 entries in total. These resources can be shared with policies that use IPv4 criteria, though it consumes two entries in hardware consumed per IPv4 criteria entry. It allows user to make use of spare IPv6 resources for IPv4 criteria policies, though if user plans to have a larger number of IPv4 criteria policies they are better off allocating more resources for use with IPv4 criteria.</p> <p>Note when a slice is allocated to IPv6 criteria, software automatically adjusts the number of available entries in that slice to 128, instead of 256, because 2 entries are needed to match IPv6 fields. The number of meters available does not reduce though and 128 meters are available for use.</p> |
| dot1p-only, IPv4 dscp-only, IPv6 dscp-only, and default SAP ingress QoS policies | <p>User can use the option 'dot1p-only' or dscp-only', if they plan to use only dot1p bits or only DSCP bits for SAP ingress classification. This typically allows for efficient use of available hardware resources and better scaling. SAP ingress policies that use only dot1p bits or only IPv4/IPv6 DSCP and Default SAP ingress QoS policies can use the resources from slices currently allocated for use by either IP-criteria or MAC-criteria or IPv6 criteria.</p> <p>The following are some special cases noted below for allocation of resources for default, dot1p-only and dscp-only SAP ingress policies:</p> <ul style="list-style-type: none"> If no slices are available to accommodate a SAP that is associated with default or dot1p-only, or a dscp-only SAP ingress policy, the software allocates resources against mac-criteria when the SAP is configured in a VLL or VPLS service. The software uses the required number of entries for this policy. The remaining entries is available for SAPs that use mac-criteria or that use only dot1p or only ipv4/ipv6 DSCP or that use default policy. |

| Type of match criteria | Description |
|------------------------|---|
| | <ul style="list-style-type: none">If no slices are available to accommodate a SAP that is associated with default, dot1p-only, or a dscp-only SAP ingress policy, the software allocates resources against ipv4-criteria when the SAP is configured in an IES or a VPRN service. The software uses the required number of entries for this policy. The remaining entries are available for SAPs that use ipv4-criteria or that use only ipv4/ipv6 DSCP or only dot1p criteria or that use the default policy. |

The SAP ingress resource slices described in this section are different from the resources that are configured using the **num-qos-classifiers** command. The **num-qos-classifiers** command sets the limit on the resources needed per SAP ingress QoS policy. The resources set the maximum limit on the resources available for use by all the SAP ingress policies that are in use simultaneously on the system. The SR OS manages the resource slices allocated to the SAP ingress QoS policy pool and allocates the entries in the slices when a SAP ingress QoS policy is associated with a SAP. This means that a SAP specifies the amount of QoS resources it needs, using the **num-qos-resources** command (in the SAP ingress policy) and the system allocates the resources required by a SAP from the slices, depending on whether the SAP ingress policy uses ip-criteria, mac-criteria, or ipv6-criteria.

**Note:**

On the 7210 SAS-D and 7210 SAS-Dxp, mac-criteria SAP ingress QoS policies can use an additional 128 classification entries with 64 meters. These entries are allocated to the mac-criteria SAP ingress QoS resource pool by default and cannot be reassigned to any another feature or any other match criteria.

Use the **tools dump system-resources** command to display information about the current usage and availability. One or more entries per slice are reserved for system use.

9.1.3 Computation of SAP ingress classification and meter resources used per SAP ingress policy

**Note:**

See the *7210 SAS Release Notes* for information about services supported on different 7210 SAS platforms. The references to services in the following table appear for completeness and do not imply support is available.

This section provides information about the resource consumption per SAP ingress policy. Resources required by the SAP ingress policy are allocated from the ingress-internal-tcam resource pool, based on the amount of resources allocated for SAP ingress classification.

The user can configure the number of classification entries the SAP requires (for example, TQ).

Number of meters allocated automatically by system = $TQ / 2$

(up to a maximum of 32 meters)

To calculate the number of SAPs allowed, assume all configured to use 'TQ' QoS resources per SAP.

Number of SAPs allowed = maximum classification entries / TQ



Note:

The number of SAPs arrived at using the preceding equation is subject to system limits. The preceding equation is used to derive the limit on the number of SAPs due to QoS resources only.

The user is allowed to mix and match SAPs with different QoS resources (that is, using different values of TQ). The allowed values in 7210 SAS-D devices for the **num-qos-classifiers** command are: 4, 8, 16, 32, 64, 128, and 256. The allowed **num-qos-classifiers** values for 7210 SAS-Dxp devices are any multiple of 2, between 2 and 256.

The following items determine the number of QoS resources to be allocated for a SAP:

- number of match-criteria entries used to identify the FC
- number of FCs to use and number of traffic-types to be policed per FC
- the amount of hardware classification resources needed per entry configured by the user (see [Resource allocation for service ingress QoS policy classification rules](#) for information about resources needed per match entry; it varies based on different match criteria in use)

Only those FCs that are in use by the match-criteria classification entries are considered for the number of FCs. These FCs are referred to as "FCs in use".

9.1.3.1 Determining the number of classification entries

This section describes the rules and methods of determining the number of classification entries.

9.1.3.1.1 Rules for a SAP in a VPLS

Given the number of traffic types to use per "FC in use", the following rules apply for a SAP in a VPLS service to determine the number of classification entries per FC in use:

- If an FC is in use and is created without explicit meters, use default meter 1 for unicast traffic and for all other traffic types (that is, broadcast, multicast, and unknown-unicast). This requires one classification entry in hardware. This assumes default multipoint meter 11 is not created by the user.
- If an FC is in use and is created without explicit meters, use default meter 1 for unicast traffic and default meter 11 (assuming meter "11" is created by the user), for all other traffic types (that is, broadcast, multicast, and unknown-unicast). This requires two classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and for all other traffic types (that is, broadcast, multicast, and unknown-unicast). This requires one classification entry in hardware. This assumes default multipoint meter "11" is not created by the user.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter 11 (assuming meter "11" is created by the user) for all other traffic types. This requires two classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively and use the unicast meter for all other traffic types (that is, multicast and unknown-unicast). This requires two classification entries in hardware. This assumes that the default multipoint meter 11 is not created by the user.
- If an FC is in use and is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively and use meter 11 (assuming meter 11 is created by the user) for all other traffic types. This requires three classification entries in hardware.

- If an FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic. This requires two classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter, an explicit broadcast meter, and an explicit multicast meter, use these meters for unicast, broadcast, and multicast traffic types respectively. The unknown unicast traffic type uses the explicitly defined multicast meter. This requires three classification entries in hardware.

9.1.3.1.2 Rules for a SAP in a VLL, IES, or VPRN service with PIM/IP multicast disabled

For calculating the number of classification entries per FC for a SAP in a VLL, IES or VPRN service with PIM/ IP multicast disabled, the following rules apply:

- Multipoint meters cannot be used. Multipoint meters configured in a policy are not used when the policy is applied to a SAP in an Epipe service.
- All FCs in use and associated with a meter always use the unicast meter. Therefore, all FCs in use only use one classification entry in the hardware.

9.1.3.1.3 Rules for a SAP in an IES or VPRN service with PIM/IP multicast enabled

Given the number of traffic types to use per "FC in use", the following rules apply for a SAP in an IES and VPRN service enabled with PIM/IP multicast enabled to arrive at a number of classification entries per FC in use:

- If an FC is in use and is created without explicit meters, use default meter 1 for unicast traffic and multicast traffic. This requires one classification entry in hardware. This assumes the default multipoint meter 11 is not created by the user.
- If an FC is in use and is created without explicit meters, use default meter 1 for unicast traffic and default meter 11 (assuming meter "11" is created by the user) for multicast traffic. This requires two classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and for multicast traffic. This requires one classification entry in hardware. This assumes default multipoint meter "11" is not created by the user.
- If an FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter 11 (assuming meter "11" is created by the user) for multicast traffic. This requires two classification entries in hardware.
- If an FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for multicast traffic. This requires two classification entries in hardware.

9.1.3.1.4 Calculating the number of classification entries per FC

Apply the rules to determine the number of classification entries per FC (only for the FCs in use) using the following equation:

$$C(i) = SFC_i(\text{unicast}) + FC_i(\text{multicast}) + FC_i(\text{broadcast}) + FC_i(\text{unknown_unicast})$$

$i = nc, h1, ef, h2, l1, af, l2, be$

where:

- FC_i (unicast), FC_i (multicast), FC_i (broadcast), and FC_i (unknown-unicast) are set to a value of 1 if this FC uses classifier to identify traffic-type unicast, multicast, broadcast, and unknown-unicast respectively
- FC_i (unicast), FC_i (multicast), FC_i (broadcast), and FC_i (unknown-unicast) are set to a value of 0 if this FC does not use a classifier to identify this traffic-type

If the user does not configure meters explicitly for the FC and meter "11" is not created, the default unicast meter is used for all traffic types; therefore only one classification entry in hardware is required by the FC. If the user does not configure meters explicitly for the FC and meter "11" is created, the default unicast meter and multicast meter are used. By default, two classification entries in hardware are required by an FC.

9.1.3.1.5 Determining the number of classification entries per policy (TC)

Taking into account the number of match criteria and the number of FCs used, use the following equation to determine the total number of classification entries per policy.

$$TC = S E(i) * C(i)$$

$i = nc, h1, ef, h2, l1, af, l2, be$

where:

- $E(i)$ is the number of match-criteria entries that classify packets to FC_i . For 7210 SAS platforms, the maximum number of classification entries per policy can be 64 (including default).
- $C(i)$ is the number of classification entries that are required by FC_i to identify different traffic types.

9.1.3.1.6 Determining the number of policers/meters per policy (TP)

Determine the number of policers or meters to use (for example TP). A maximum of 32 meters per policy are available.

Only those meters associated with FCs are considered for number of meters. Note that only "FCs in use" is considered.

$$\text{Total QoS resources required (TQ)} = \max((TC), (2 * TP))$$

The number obtained is rounded off to next multiple of "2" greater than the TQ obtained using this equation.

The user configures the value of TQ using the **num-qos-classifiers** command.

For more examples and information about resource allocation, see [Service ingress QoS policy configuration considerations](#).

9.1.3.2 Service ingress QoS policies resource usage examples

This section provides resource usage examples for service ingress QoS policies.

9.1.3.2.1 Example 1

```
sap-ingress 10 create
description "example-policy-1"
  num-qos-classifiers 8
  meter 1 create
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 11 multipoint create
  exit
  fc "af" create
    meter 1
  exit
  fc "be" create
    meter 3
  exit
  fc "h2" create
    meter 3
  exit
  fc "l1" create
    meter 3
  exit
  mac-criteria
    entry 1 create
      match
        dot1p 7 7
      exit
      action fc "af"
    exit
    entry 2 create
      match
        dot1p 5 7
      exit
      action fc "l1"
    exit
    entry 3 create
      match
        dot1p 6 7
      exit
      action fc "h2"
    exit
  default-fc "be"
```

In example 1, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC use the following computation.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 1 + 0 = 2$$

Because this FC uses a unicast meter, an entry is needed to identify this traffic type explicitly. Another entry is needed to classify broadcast, multicast, and unknown-unicast traffic type to the same FC and use the default meter 11.

$$FCI1 = 1 + 0 + 1 + 0 = 2$$

$$FCaf = 1 + 0 + 1 + 0 = 2$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 1 + 0 = 2$$

Use the following equation to calculate the total classification entries used by this policy.

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (1 * 2)h2 + (1 * 2)l1 + (1 * 2)af + (0 * 0)l2 + (1 * 2)be = 8$$

(three explicit match criteria entries are used to map traffic to each of FC H2, FC L1, and FC AF along with a default classification entry for FC BE).

The number of meters used = 3 (because both FCs use meter 1, meter 3, and meter 11).

In this example, **num-qos-classifiers 16** is used (that is, maximum of (8, (2 * 3))).

If the same policy were used for a SAP in an Epipe service, and all traffic is classified to a unicast traffic type and only unicast meters are used, the following computation is used.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the preceding equation, calculate the total classification entries used = 4 and meters used = 2.

9.1.3.2.2 Example 1a (default multipoint meter 11 is not used)

```
sap-ingress 10 create
description "example-policy"
num-qos-classifiers 4
meter 1 create
rate cir 0 pir max
exit
meter 3 create
```



```

rate cir 100 pir 100
exit

scope template

default-fc be

fc be create
meter 3
exit
fc af create
meter 1
exit
fc l1 create
meter 3
exit
fc h2 create
meter 3
exit
mac-criteria dot1p-only
entry 1 create
match dot1p 7
action fc af
exit
entry 2 create
match dot1p 5
action fc l1
exit
entry 3 create
match dot1p 6
action fc h2
exit
exit
exit

```

In example 1a, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC use the following computation.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

Because this FC uses unicast meter for all traffic types, an entry to classify all traffic types to this FC explicitly is required.

$$FCI1 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Calculate the total classification entries used by this policy, using the following equation.

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (1 * 1)h2 + (1 * 1)l1 + (1 * 1)af + (0 * 0)l2 + (1 * 1)be = 4$$

(three explicit match criteria entries are used to map traffic to each of FC H2, FC L1, and FC AF along with a default classification entry for FC BE).

The total number of meters used = 2 (because FCs use meter 1 and meter 3).

In this example, **num-qos-classifiers 4** is used (maximum of (4, (2 * 2))). The use of a unicast meter for all traffic-types allows QoS resources to be used efficiently.

If the same policy were used for a SAP in an Epipe service, and all traffic is classified to a unicast traffic type and only unicast meters are used, the following computations are used.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCl1 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCl2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using this equation, total classification entries used = 4 and meters used = 2.

Using the same policy for Epipe SAPs can lead to inefficient use of resources. Nokia recommends creating a different policy with the required number of resources (that is, with **num-qos-classifiers 4**).

9.1.3.2.3 Example 2

```
sap-ingress 10 create
description"example-policy-1"
num-qos-classifiers16
    meter 1 create
    exit
    meter 2 create
    rate cir 1 pir 20
    exit
    meter 3 create
    rate cir 100 pir 100
    exit
    meter 11 multipoint create
    exit
    fc "af" create
    meter 3
    broadcast-meter 2
    exit
    fc "be" create
    meter 3
    broadcast-meter 2
```

```
exit
fc "h2" create
    meter 3
    broadcast-meter 2
exit
fc "l1" create
    meter 3
    broadcast-meter 2
exit
mac-criteria
    entry 1 create
        match
            dot1p 7 7
        exit
        action fc "af"
    exit
    entry 2 create
        match
            dot1p 5 7
        exit
        action fc "l1"
    exit
    entry 3 create
        match
            dot1p 6 7
        exit
        action fc "h2"
    exit
exit
default-fc "be"
```

In example 2, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC use the following computation.

$$\text{FCnc} = 0 + 0 + 0 + 0 = 0$$

$$\text{FCh1} = 0 + 0 + 0 + 0 = 0$$

$$\text{FCef} = 0 + 0 + 0 + 0 = 0$$

$$\text{FCh2} = 1 + 1 + 1 + 0 = 3$$

FC uses unicast meter and broadcast meter, so two entries are required to identify these traffic types explicitly. Another entry is required to classify multicast and unknown-unicast traffic type to the same FC and use the default meter 11.

$$\text{FCI1} = 1 + 1 + 1 + 0 = 3$$

$$\text{FCaf} = 1 + 1 + 1 + 0 = 3$$

$$\text{FCI2} = 0 + 0 + 0 + 0 = 0$$

$$\text{FCbe} = 1 + 1 + 1 + 0 = 2$$

Using this equation, the total classification entries used = 11 (three explicit match criteria entries map to each of FC H2, L1, and AF along with a default classification rule for BE).

Meters used = 3 (FCs use only meter 2, meter 3, and meter 11).

Therefore, in this example, **num-qos-classifiers 16** is used (that is, a maximum of (12, (2*3))). Note that the system internally uses 18, instead of 16.

If the same policy were used for a SAP in an Epipe service, all traffic is classified to a unicast traffic type, and only unicast meters are used, use the following computation.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using this equation, the total classification entries used = 4 and the total meters used = 1.

9.1.3.2.4 Example 2a (default multipoint meter 11 is not used)

```
sap-ingress 10 create
description "example-policy-1"
num-qos-classifiers 8

meter 1 create
rate cir 0 pir max
exit
meter 3 create
rate cir 100 pir 100
exit
meter 2 create
rate cir 1 pir 20
exit
scope template
default-fc be
fc be create
meter 3
broadcast-meter 2
exit
fc af create
meter 3
broadcast-meter 2
exit
fc l1 create
meter 3
broadcast-meter 2
exit
fc h2 create
meter 3
broadcast-meter 2
exit
```

```
mac-criteria dot1p-only
entry 1 create
match dot1p 7
action fc af
exit
entry 2 create
match dot1p 5
action fc l1
exit
entry 3 create
match dot1p 6
action fc h2
exit
exit
```

In example 2a, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC use the following computation.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 1 + 0 = 2$$

Because this FC uses a unicast meter for unicast, multicast, unknown-unicast traffic, and broadcast meter for broadcast traffic, so two entries are needed.

$$FCI1 = 1 + 0 + 1 + 0 = 2$$

$$FCaf = 1 + 0 + 1 + 0 = 2$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 1 + 0 = 2$$

Using this equation to get the total classification, entries used = 8 (three explicit match criteria entries map to each of FC H2, L1, and AF along with a default classification rule for BE).

The number of meters used = 2 (FCs use only meter 2 and meter 3).

Hence, in this example **num-qos-classifiers 8** is used (that is, maximum of (8, (2*2))).

If the same policy were used for a SAP in an Epipe service, all traffic is classified to a unicast traffic type and only unicast meters are used, use the following computation.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the preceding equation, to get total classification entries used = 4 and meters used = 1. As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Nokia recommends creating a different policy with the required number of resources (that is, with **num-qos-classifiers 4**).

9.1.3.2.5 Example 3

```
sap-ingress 10 create
description"example-policy-2"
num-qos-classifiers 16
    meter 1 create
        rate cir 100 pir 100
    exit
    meter 2 create
        rate cir 1 pir 20
    exit
    meter 3 create
        rate cir 100 pir 100
    exit
    meter 4 create
        rate cir 10 pir 100
    exit
    meter 5 create
        rate cir 10 pir 10
    exit
    meter 11 multipoint create
        rate cir 1 pir 20
    exit
    fc "af" create
        meter 3
        broadcast-meter 2
        multicast-meter 4
    exit
    fc "h1" create
        meter 5
        broadcast-meter 4
        multicast-meter 4
        unknown-meter 4
    exit
    fc "h2" create
        meter 3
        broadcast-meter 2
    exit
    fc "l1" create
        meter 3
        broadcast-meter 2
    exit
    mac-criteria
        entry 1 create
            match
                dot1p 7 7
            exit
            action fc "af"
```

```
exit
entry 2 create
  match
    dot1p 5 7
  exit
  action fc "l1"
exit
entry 3 create
  match
    dot1p 6 7
  exit
  action fc "h2"
exit
entry 4 create
  match
    dot1p 3 7
  exit
  action fc "h1"
exit
exit
default-fc "be"
```

In example 3, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC use the following computation.

$$\text{FCnc} = 0 + 0 + 0 + 0 = 0$$

$$\text{FCh1} = 1 + 1 + 1 + 1 = 4$$

Because this FC uses unicast, broadcast, multicast, and unknown-unicast meter, four entries are required to identify these traffic types explicitly.

$$\text{FCef} = 0 + 0 + 0 + 0 = 0$$

$$\text{FCh2} = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast meter and broadcast meter, two entries are required to identify these traffic types explicitly. Another entry is required to classify multicast and unknown-unicast traffic type to the same FC and use the default meter 11.

$$\text{FCI1} = 1 + 1 + 1 + 0 = 3$$

Because this FC uses only unicast meter, an entry is required to identify this traffic type explicitly. Another entry is required to classify broadcast, multicast, and unknown-unicast traffic type to the same FC and use the default meter 11.

$$\text{FCaf} = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast, broadcast, and multicast meter, three entries are required to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$\text{FCI2} = 0 + 0 + 0 + 0 = 0$$

$$\text{FCbe} = 1 + 0 + 1 + 0 = 2$$

Using the preceding equation, the total classification entries used = 15 and the total meters used = 6.

If the same policy were used for a SAP in an Epipe service, all traffic is classified to a unicast traffic type and only unicast meters are used, use the following computation.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 1 + 0 + 0 + 0 = 1$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCl1 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the preceding equation, the total classification entries used = 5 and the total meters used = 3 (because all FCs used only meter 1, meter 3, and meter 5).

9.1.3.2.6 Example 3a (default multipoint meter 11 is not used)

```
sap-ingress 10 create
description "example-policy-2"
num-qos-classifiers 12

meter 1 create
rate cir 100 pir 100
exit
meter 3 create
rate cir 100 pir 100
exit
meter 2 create
rate cir 1 pir 20
exit
meter 4 create
rate cir 10 pir 100
exit
meter 5 create
rate cir 10 pir 10
exit
scope template
default-fc be
fc af create
meter 3
broadcast-meter 2
multicast-meter 4
exit
fc l1 create
meter 3
broadcast-meter 2
exit
fc h2 create
```



```
meter 3
broadcast-meter 2
exit
fc h1 create
meter 5
broadcast-meter 4
multicast-meter 4
unknown-meter 4
exit
mac-criteria dot1p-only
entry 1 create
match dot1p 7
action fc af
exit
entry 2 create
match dot1p 5
action fc l1
exit
entry 3 create
match dot1p 6
action fc h2
exit
entry 4 create
match dot1p 3
action fc h1
exit
exit
```

In example 3a, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC use the following computation.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 1 + 1 + 1 + 1 = 4$$

Because this FC uses unicast, broadcast, multicast, and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 1 + 0 = 2$$

Because this FC uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly, multicast and unknown-unicast traffic use the same resource as the unicast traffic.

$$FCI1 = 1 + 0 + 1 + 0 = 2$$

Because this FC uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Multicast and unknown-unicast traffic use the same resource as unicast traffic.

$$FCaf = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast, broadcast, and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Because no explicit meters are configured for FC "be", it uses meter 1 for all traffic types and needs one entry to identify these traffic types.

Using the preceding equation, the total classification entries used = 12 and meters used = 5. The **num-qos-classifiers** CLI command can be set to 12 (the minimum value).

Using the same policy for a SAP in an Epipe service, all traffic is classified to a unicast traffic type and only unicast meters are used, the following computations are used.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 1 + 0 + 0 + 0 = 1$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the preceding equation, the total classification entries used = 5 and meters used = 3 (because all FCs used only meter 1, meter 3, and meter 5). For Epipe service, a policy with **num-qos-resources 6** can be used.

9.1.3.2.7 Example 4

```
sap-ingress 10 create
description"example-policy-3"
num-qos-classifiers 36
    meter 1 create
        rate cir 100 pir 100
    exit
    meter 2 create
        rate cir 1 pir 20
    exit
    meter 3 create
        rate cir 100 pir 100
    exit
    meter 4 create
        rate cir 10 pir 100
    exit
    meter 5 create
        rate cir 10 pir 10
    exit
    meter 6 create
        rate cir 11 pir 100
    exit
    meter 8 create
        rate cir 20 pir 100
```

```
exit
meter 11 multipoint create
    rate cir 1 pir 20
exit
fc "af" create
    meter 3
    broadcast-meter 2
    multicast-meter 4
exit
fc "ef" create
    meter 6
    broadcast-meter 2
    multicast-meter 8
exit
fc "h1" create
    meter 5
    broadcast-meter 4
    multicast-meter 4
    unknown-meter 4
exit
fc "h2" create
    meter 3
    broadcast-meter 2
exit
fc "l1" create
    meter 3
    broadcast-meter 2
exit
fc "nc" create
    meter 6
    broadcast-meter 2
    multicast-meter 8
exit
mac-criteria
    entry 1 create
        match
            dot1p 4 7
        exit
        action fc "af"
    exit
    entry 2 create
        match
            dot1p 5 7
        exit
        action fc "l1"
    exit
    entry 3 create
        match
            dot1p 6 7
        exit
        action fc "h2"
    exit
    entry 4 create
        match
            dot1p 3 7
        exit
        action fc "h1"
    exit
    entry 5 create
        match
            dot1p 2 7
        exit
        action fc "ef"
    exit
```

```

        entry 6 create
            match
                dot1p 7 7
            exit
            action fc "nc"
        exit
    exit
    default-fc "be"

```

In example 4, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC use the following computation.

$$FCnc = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast, broadcast, and multicast meter, three entries are required to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh1 = 1 + 1 + 1 + 1 = 4$$

Because this FC uses unicast, broadcast, multicast and unknown-unicast meter, four entries are required to identify these traffic types explicitly.

$$FCef = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast, broadcast, and multicast meter, three entries are required to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh2 = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast meter and broadcast meter, two entries are required to identify these traffic types explicitly. Another entry is required to classify multicast and unknown-unicast traffic type to the same FC and use the default meter 11.

$$FCI1 = 1 + 1 + 1 + 0 = 3$$

$$FCaf = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast, broadcast, and multicast meter, three entries are required to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 1 + 0 = 2$$

Using the preceding equation, the total classification entries used = 21 and the total meters used = 8.

If the same policy were used for a SAP in an Epipe service, all traffic is classified to a unicast traffic type and only unicast meters are used, use the following computation.

$$FCnc = 1 + 0 + 0 + 0 = 1$$

$$FCh1 = 1 + 0 + 0 + 0 = 1$$

$$FCef = 1 + 0 + 0 + 0 = 1$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the preceding equation, the total classification entries used = 7 and the total meters used = 4.

Using the same policy for Epipe SAP can lead to inefficient use of resources. Nokia recommends creating a different policy with the required number of resources (that is, **num-qos-classifiers 16**)

9.1.3.2.8 Example 4a (default multipoint meter 11 is not used)

```
sap-ingress 10 create
description "example-policy-3"
num-qos-classifiers 20
meter 1 create
rate cir 100 pir 100
exit
meter 3 create
rate cir 100 pir 100
exit
meter 2 create
rate cir 1 pir 20
exit
meter 4 create
rate cir 10 pir 100
exit
meter 5 create
rate cir 10 pir 10
exit
meter 6 create
rate cir 11 pir 100
exit
meter 8 create
rate cir 20 pir 100
exit
scope template
default-fc be
fc af create
meter 3
broadcast-meter 2
multicast-meter 4
exit
fc l1 create
meter 3
broadcast-meter 2
exit
fc h2 create
meter 3
```

```
broadcast-meter 2
exit
fc h1 create
meter 5
broadcast-meter 4
multicast-meter 4
unknown-meter 4
exit
fc ef create
meter 6
broadcast-meter 2
multicast-meter 8
exit
fc nc create
meter 6
broadcast-meter 2
multicast-meter 8
exit
mac-criteria dot1p-only
entry 1 create
match dot1p 4
action fc af
exit
entry 2 create
match dot1p 5
action fc l1
exit
entry 3 create
match dot1p 6
action fc h2
exit
entry 4 create
match dot1p 3
action fc h1
exit
entry 5 create
match dot1p 2
action fc ef
exit
entry 6 create
match dot1p 7
action fc nc
exit
exit
exit
```

In example 4a, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC use the following computation.

$$FCnc = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast, broadcast, and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh1 = 1 + 1 + 1 + 1 = 4$$

Because this FC uses unicast, broadcast, multicast, and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$FCef = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast, broadcast, and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh2 = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Multicast and unknown-unicast traffic of the same FC use the unicast resources (both meter and classification entry).

$$FCI1 = 1 + 1 + 1 + 0 = 3$$

$$FCaf = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast, broadcast, and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Because this FC uses a single meter for all traffic-types only, a single meter and single entry is needed.

Using the preceding equation, the total classification entries used = 20 and meters used = 7. Therefore, in this example **num-qos-classifiers 20** is used (the minimum value).

If the same policy were used for a SAP in an Epipe service, all traffic is classified to a unicast traffic type, and only unicast meters are used, use the following computation.

$$FCnc = 1 + 0 + 0 + 0 = 1$$

$$FCh1 = 1 + 0 + 0 + 0 = 1$$

$$FCef = 1 + 0 + 0 + 0 = 1$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FCI1 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the preceding equation, the total classification entries used = 7 and meters used = 4.

Using the same policy for Epipe SAP can lead to inefficient use of resources. Nokia recommends creating a different policy with the required number of resources (that is, with **num-qos-classifiers 8**).

9.1.3.2.9 Example 5

```
sap-ingress 10 create

  num-qos-classifiers 72
  meter 1 create
  exit
  meter 3 create
  exit
  meter 4 create
  exit
  meter 11 multipoint create
  exit
  fc "af" create
    meter 3
    broadcast-meter 11
    multicast-meter 4
  exit
  fc "be" create
    meter 1
    broadcast-meter 11
  exit
  ip-criteria
    entry 1 create
      match
        dscp be
      exit
      action fc "af"
    exit
    entry 2 create
      match
        dscp cp1
      exit
      action fc "af"
    exit
    entry 3 create
      match
        dscp cp3
      exit
      action fc "af"
    exit
    entry 4 create
      match
        dscp cp4
      exit
      action fc "af"
    exit
    entry 5 create
      match
        dscp cp5
      exit
      action fc "af"
    exit
    entry 6 create
      match
        dscp cp6
      exit
      action fc "af"
    exit
    entry 7 create
      match
        dscp cp7
```



```

        exit
        action fc "af"
    exit
    entry 8 create
        match
            dscp cs1
        exit
        action fc "af"
    exit
    entry 9 create
        match
            dscp cp9
        exit
        action fc "af"
    exit
    entry 10 create
        match
            dscp af11
        exit
        action fc "af"
    exit
    entry 11 create
        match
            dscp cp11
        exit
        action fc "af"
    exit
    entry 12 create
        match
            dscp af12
        exit
        action fc "af"
    exit
    entry 13 create
        match
            dscp cp13
        exit
        action fc "af"
    exit
    entry 14 create
        match
            dscp cp15
        exit
        action fc "af"
    exit
    entry 15 create
        match
            dscp cp15
        exit
        action fc "af"
    exit
    exit
    default-fc "be"

```

In example 5, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC use the following computation.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 0 + 0 + 0 + 0 = 0$$

$$FCI1 = 0 + 0 + 0 + 0 = 0$$

$$FCaf = 1 + 0 + 1 + 0 = 3$$

Because this FC uses unicast meter, an entry is required to identify these traffic types explicitly. Another entry is required to classify broadcast, multicast, and unknown-unicast traffic type to the same FC and use the default meter 11.

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 1 + 1 + 0 = 3$$

Because this FC uses unicast, broadcast, and multicast meter, three entries are required to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

Calculate the total classification entries used by this policy using the following equation.

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (15 * 3)af + (0 * 0)l2 + (1 * 3)be = 48$$

The total meters used in this policy = 4.

In example 5, **num-qos-classifiers 72** is used (that is, a maximum of $(48, (2 * 4)) = 48$, rounded off to the next available numQosClassifier range.

If the same policy were used for a SAP in an Epipe service, all traffic is classified to a unicast traffic type and only unicast meters are used, use the following computation.

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 0 + 0 + 0 + 0 = 0$$

$$FCI1 = 0 + 0 + 0 + 0 = 0$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FCI2 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Calculate the total classification entries used by this policy using the following equation.

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (15 * 1)af + (0 * 0)l2 + (1 * 1)be = 16$$

The total meters used in this policy = 2.

For Epipe SAP, Nokia recommends defining another sap-ingress policy using **num-qos-classifiers 16** (maximum of $(16, (2 * 2)) = 16$).

9.2 Basic configurations

A basic service ingress QoS policy must conform to the following:

- have a unique service ingress QoS policy ID
- allocate the number of classifier and meter resources needed for use
- have a QoS policy scope of **template** or **exclusive**
- have at least one default unicast FC meter/queue
- (optional) use multipoint FC meter/queue

9.2.1 Create service ingress QoS policies

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP, a default QoS policy is applied.

9.2.1.1 Service ingress QoS policy



Note:

A meter is available to limit the bandwidth per FC on service ingress.

To create a service ingress policy, perform the following.

1. Define a policy ID value. The system does not dynamically assign a value.
2. Define a description that provides a brief overview of policy features.
3. Specify the **num-qos-classifiers** command. The default value is 2. The number of meters/queue allocated is equal to half the number of classifiers specified.
4. Specify a default FC for the policy. All packets received on an ingress SAP using this ingress QoS policy will be classified to the default FC.
5. Define FC parameters by performing the following:
 - Modify the unicast-meter/queue default value to override the default unicast forwarding type meter mapping for **fc fc-name**.
 - Modify the **multicast-meter/queue** default value to override the default multicast forwarding type meters/queue mapping for **fc fc-name**.
 - Modify the **multicast-meter/queue** default value to override the default unknown unicast forwarding type **meter** mapping for **fc fc-name**.
 - Modify the **broadcast-meter** default value to override the default broadcast forwarding type **meter** mapping for **fc fc-name**.
6. On platforms where applicable, 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 ingress meter and corresponding FC for matched traffic.

A SAP ingress policy is created with a **template** scope. The scope can be modified to **exclusive** for a special one-time use policy. Otherwise, the **template** scope enables the policy to be applied to multiple SAPs.

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#
```

9.2.1.1.1 Service ingress QoS meter

About this task

To create service ingress meter parameters, perform the following.

Procedure

- Step 1.** Define a new meter ID value; the system does not dynamically assign a value.
- Step 2.** Configure meter parameters. Ingress meters support the definition of either srTCM (Single Rate Tri-Color Meter) or trTCM (Two Rate Tri-Color Meter), CIR/PIR, CBS/MBS parameters.

Example

The following is a sample ingress meter configuration output.

```
A:ALA-7>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
...
sap-ingress 100 create
description "Used on VPN sap"
meter 1 create
exit
meter 11 multipoint create
exit
meter 2 create
rate cir 11000
exit
meter 3 create
cbs 32
rate cir 11000
exit
meter 4 create
rate cir 100 pir 500
exit
meter 5 create
cbs 64
mbs 128
rate cir 1500 pir 1500
exit
meter 6 create
mode srtcm
rate cir 2500 pir 2500
exit
meter 7 create
```

```
cbs 256
mbs 512
rate cir 100 pir 36
exit
meter 8 create
cbs 256
mbs 512
rate cir 11000
exit
meter 9 create
rate cir 11000
exit
meter 10 create
rate cir 1
exit
meter 12 create
rate cir 1500 pir 1500
exit
meter 13 create
rate cir 2500 pir 2500
exit
meter 14 create
rate cir 36 pir 100
exit
meter 15 create
rate cir 36 pir 100
exit
meter 16 create
cbs 128
mbs 256
rate cir 36 pir 100
exit
...
#-----
A:ALA-7>config>qos#
```

9.2.1.1.2 SAP ingress FC

Example

The following is a sample FC configuration output.

```
A:ALA-7>config>qos# info
#-----
...
fc af create
meter 1
broadcast-meter 7
unknown-meter 8
exit
fc be create
meter 2
unknown-meter 9
exit
fc ef create
meter 3
broadcast-meter 10
exit
fc h1 create
meter 4
```

```
multicast-meter 12
exit
fc h2 create
meter 5
broadcast-meter 13
multicast-meter 14
unknown-meter 15
exit
fc nc create
meter 6
broadcast-meter 16
multicast-meter 17
unknown-meter 18
exit

...
#-----
```

9.2.1.1.3 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:

The following is a sample 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#
```

9.2.1.1.4 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#
```

9.2.1.2 Applying service ingress policies

This section describes applying SAP ingress policies to service SAPs.

9.2.1.2.1 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
-----
A:ALA-7>config>service#

```

9.2.1.2.2 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#

```

9.2.1.2.3 IES



Note:

SAP ingress QoS policies for access SAPs and IES on access SAPs are only supported on 7210 SAS-D and 7210 SAS-Dxp.

Example

The following is a sample IES service configuration output.

```

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
...

```



```
-----
A:ALA-7>config>service#
```

9.3 Service management tasks

This section describes the 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.

```
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

The following Epipe service output examples show that the SAP service ingress reverts to *policy-id* 1 when the non-default policies are removed from the configuration.

Example

```
A:ALA-104>config>service>epipe# info detail
-----
description "Distributed Epipe service to west coast"
    no tod-suite
    dotlag
    exit
    ingress
        qos 1
        no filter
    exit
    egress
        no filter
    exit
    no collect-stats
    no accounting-policy
    no shutdown
-----
A:ALA-7>config>service>epipe#
```

9.3.2 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]
```

9.3.3 Remove a policy from the QoS configuration

Use the following syntax to remove a configuration 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

- [Service ingress QoS policy commands](#)
- [Operational commands](#)
- [Show commands](#)

9.4.1 Service ingress QoS policy commands

```
- config
  - qos
    - [no] sap-ingress policy-id [create]
      - default-fc fc
      - no default-fc
      - description description-string
      - no description
      - [no] fc fc-name [create]
        - broadcast-meter meter-id
        - no broadcast-meter
        - meter meter-id
        - no meter
        - multicast-meter meter-id
        - no multicast-meter
```

```

- unknown-meter meter-id
- no unknown-meter
- [no] ip-mac-match {ip-first | mac-first}
- [no] ip-criteria [any | dscp-only]
  - [no] entry entry-id [create]
    - action [fc fc-name]
    - no action
    - description description-string
    - no description
    - match [protocol protocol-id]
    - no match
      - dscp dscp-value | dscp-name [dscp-mask]
      - no dscp
      - dst-ip {ip-address/mask | ip-address netmask}
      - no dst-ip
      - dst-port {eq} dst-port-number
      - no dst-port
      - ip-prec ip-prec-value [ip-prec-mask]
      - no ip-prec
      - src-ip ip-address/mask
      - no src-ip
      - src-port {eq} src-port-number
      - no src-port
    - renum [old-entry-id new-entry-id]
- [no] ipv6-criteria [any | dscp-only] [IPv6 Match Criteria]
  - [no] entry entry-id [create]
    - action [fc fc-name]
    - no action
    - description description-string
    - no description
    - match [next-header next-header]
    - no match
      - dscp dscp-value | dscp-name [dscp-mask]
      - no dscp
      - dst-ip {ipv6-address/prefix-length}
      - no dst-ip
      - dst-port {eq} dst-port-number
      - no dst-port
      - ip-prec ip-prec-value [ip-prec-mask]
      - no ip-prec
      - src-ip {ipv6-address/prefix-length}
      - no src-ip
      - src-port {eq} src-port-number
      - no src-port
    - renum [old-entry-id new-entry-id]
- [no] mac-criteria [any | dot1p-only]
  - [no] entry entry-id [create]
    - action [fc fc-name]
    - no action
    - description description-string
    - no description
    - [no] match
      - dot1p dot1p-value [dot1p-mask]
      - no dot1p
      - dst-mac ieee-address [ieee-address-mask]
      - no dst-mac
      - etype 0x0600..0xffff
      - no etype
      - src-mac ieee-address [ieee-address-mask]
      - no src-mac
    - renum old-entry-id new-entry-id
- meter meter-id [multipoint] [create]
- no meter meter-id
  - adaptation-rule [cir adaptation-rule] [pir adaptation-rule]

```

```
- no adaptation-rule
- cbs size [kbits | bytes | kbytes]
- no cbs
- color-mode color-mode
- no color-mode
- mbs size [kbits | bytes | kbytes]
- no mbs
- mode mode
- no mode
- rate cir-rate-in-kbps [pir pir-rate-in-kbps]
- no rate
- num-qos-classifiers [num-resources] [ipv6 | no-ipv6]
- scope {exclusive | template}
- no scope
```

9.4.2 Operational commands

```
- config
  - qos
    - copy sap-ingress src-pol dst-pol [overwrite]
```

9.4.3 Show commands

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

9.5 Command descriptions

9.5.1 Configuration commands

9.5.1.1 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.

9.5.1.2 Operational commands



Note:

QoS capabilities vary across 7210 SAS platforms. The terms meter/queue or queue/meter are used in the command descriptions; meters, queues, or both may apply depending on the capabilities of the 7210 SAS platform. The descriptions of specific commands also mention the capabilities of the platform/node, where applicable.

copy

Syntax

```
copy sap-ingress src-pol dst-pol [overwrite]
```

Context

```
config>qos
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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 based on 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 by the contents of the source policy. If **overwrite** is not specified, an error occurs if the destination policy ID exists.

renum

Syntax

renum *old-entry-id new-entry-id*

Context

config>qos>sap-ingress>ip-criteria

config>qos>sap-ingress>ipv6-criteria

config>qos>sap-ingress>mac-criteria

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command renumbers existing QoS policy criteria entries to properly sequence policy entries.

This can be required in some cases because the 7210 SAS exits when the first match is found and executes the actions in accordance with the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.

Parameters

old-entry-id

Specifies the entry 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

9.5.1.3 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 is used to create or edit the ingress policy. The ingress policy defines the Service Level Agreement (SLA) enforcement service packets receive when they ingress a SAP. SLA enforcement requires the definition of meters/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/ queues 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, IP DSCP, IP headers, MAC headers, or all as the match 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.

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. The work area policy can be modified until complete and then written over the original policy ID. Use the **config qos copy** command to maintain policies in this manner.



Note:

- Before a SAP ingress policy can be associated with a SAP, resources must be allocated using the **config>system>resource-profile>ingress-internal-tcam>qos-sap-ingress-resource** command. See [Resource allocation for service ingress QoS policy classification rules](#) for information about resource allocation. See the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K

3SFP+ 8C Basic System Configuration Guide for more information about system resource allocation and examples for this CLI command.

- On the 7210 SAS-D and 7210 SAS-Dxp, only meters are supported on service ingress for rate enforcement. These platforms do not support service ingress queues.
- On the 7210 SAS-D and 7210 SAS-Dxp, the user has the option to use dot1p, IP DSCP, IPv4 and IPv6 criteria, and MAC criteria.

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 ID1.

Default

policy-id 1

Parameters

policy-id

Specifies a policy ID that uniquely identifies the policy.

Values 1 to 65535

create

Specifies to create a SAP ingress policy.

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.

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 SAP. If a policy with an **exclusive** scope is assigned to a second SAP, an error message is generated. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.

template

Specifies that the policy can be applied to multiple SAPs on the router. An error is generated when **scope template** is changed to **scope exclusive** for default policies.

default-fc

Syntax

default-fc *fc*

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the default FC for the policy. If an ingress packet does not match a higher priority (more explicit) classification command, the default FC is associated with the packet. Unless overridden by an explicit FC classification rule, all packets received on an ingress SAP using this ingress QoS policy is classified to the default FC.

The default FC is best effort (be). The **default-fc** settings are displayed in the **show** configuration and save output, regardless of whether the **detail** keyword is included in the **show** command.

The **in** and **out** parameters are mutually exclusive following the profile classification keyword. Either **in** or **out** must be specified when the profile keyword is present.

Default

be

Parameters

fc

Specifies the FC name for the queue/meter. The value specified for *fc* must be one of the predefined FCs in the system.

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

The **fc** command creates a class instance of the FC *fc-name*. After the *fc-name* is created, classification actions can be applied and then used in match classification criteria.

The **no** form of this command removes all explicit queue mappings for *fc-name* forwarding types. The queue mappings revert to the default meters for *fc-name*.

Parameters

fc-name

Specifies the FC name for the queue. The value specified for the *fc-name* must be one of the predefined FCs for the system.

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

create

Keyword to create a class instance of the FC.

ip-mac-match

Syntax

[no] **ip-mac-match** {**ip-first** | **mac-first**}

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the user to match on both IP and MAC criteria in a SAP ingress policy. If this command is not executed, the software does not allow for configuration of both IP and MAC criteria in a SAP ingress policy. In other words, without this command in a SAP ingress policy, IP and MAC criteria are mutually exclusive.

The user also has the option to specify if all the IP criteria entries configured in the policy need to be matched first followed by all the MAC criteria entries or the other way around. For example, if **ip-first** is configured, all the IP criteria entries are compared for matches first; if there are no matches, then MAC criteria entries are compared for matches. If a match is found, no further matches are performed and the actions associated with the matched entry are taken.

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

Default

no ip-mac-match

Parameters

ip-first

Specifies to match all the IP criteria entries first before matching any of the MAC entries.

mac-first

Specifies to match all the MAC criteria entries first before matching any of the IP entries.

ip-criteria

Syntax

[no] ip-criteria [any | dscp-only] *policy id*

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command selects the ingress meter and corresponding FC for matched traffic.

The user can specify either **any** or **dscp-only** as the sub-criteria. The sub-criteria determines what fields can be used to match traffic. The resource allocation for classification is affected by the sub-criteria in use. See [Resource allocation for service ingress QoS policy classification rules](#) for more information.

This command is used to enter the context to create or edit policy entries that specify IP criteria DiffServ code point.

The 7210 SAS implementation exits on the first match found and executes the actions in accordance with the accompanying **action** command. For this reason, entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. After an IP criteria entry is removed from a SAP ingress policy, the IP criteria is removed from all services where that policy is applied.

Default

dscp-only

Parameters

any

Specifies that entries can use any of the fields available under ip-criteria (Example - IP source, IP destination, IP protocol fields can be used) for matching

dscp-only

Specifies that entries can use the IP DSCP field or IP precedence field.

policy-id

Specifies the policy ID.

Values 1 to 65535

ipv6-criteria

Syntax

[no] **ipv6-criteria** [**any** | **dscp-only**] *policy-id*

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command selects the appropriate ingress meters and corresponding FC for matched traffic.

This command is used to enter the node to create or edit policy entries that specify IPv6 criteria, such as IP quintuple lookup or DSCP.

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 IPv6 criteria with a SAP, resources must be allocated using the **config system resource-profile ingress-internal-tcam qos-sap-ingress-resource ipv6-ipv4-match-enable** command. See the *7210 SAS-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. After an IPv6 criteria entry is removed from a SAP ingress policy, the IPv6 criteria is removed from all services where that policy is applied.

Parameters

any

Specifies that entries can use any of the fields available under ipv6-criteria (Example - IPv6 source, IPv6 destination, IPv6 protocol fields can be used) for matching

dscp-only

Specifies that entries can use the IP DSCP field or IPv6 precedence field.

policy-id

Specifies the policy ID.

Values 1 to 65535

mac-criteria

Syntax

[no] mac-criteria [any | dot1p-only] policy id

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command selects the appropriate ingress meters and corresponding FC for matched traffic.

User can specify either **any** or **dot1p-only** as the sub-criteria. The sub-criteria determines what fields can be used to match traffic. The resource allocation for classification is affected by the sub-criteria in use. See [Resource allocation for service ingress QoS policy classification rules](#) for more information about SAP ingress resource allocation.

This command is used to enter the node to create or edit policy entries that specify MAC criteria.

The 7210 SAS OS implementation exits on the first match found and executes the actions in accordance with the accompanying **action** command. For this reason, entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. After a MAC criteria entry is removed from a SAP ingress policy, the MAC criteria is removed from all services where that policy is applied.

Default

dot1p-only

Parameters

any

Specifies that entries can use the other MAC header fields for matching.

dot1p-only

Specifies that entries can use only the dot1p field.

policy-id

Specifies the policy ID.

Values 1 to 65535

num-qos-classifiers

Syntax

num-qos-classifiers [*num-resources*] [**ipv6** | **no-ipv6**]

Context

config>qos>sap-ingress>num-qos-classifiers

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the number of classifiers the SAP ingress Qos policy can use. This parameter cannot be modified when it is associated with a SAP.

The *num-resources* parameter also determines the maximum number of meters that are available to this policy. The maximum number of meters available for use by the FCs (FC) defined under this policy is equal to half the value specified in *num-resources* parameter (maximum of 32). Any of these meters are available for use to police unicast or multipoint traffic. Any of these meters is available for use by more than one FC (or a single meter is available for use by all the FCs).

The **ipv6** keyword indicates a plan to use **ipv6-criteria**, and the resources needed for this SAP ingress QoS policy must be allocated to the slice assigned to IPv6 criteria.

Default

num-qos-classifiers 2 no-ipv6

Parameters

num-resources

Specifies the number of resources planned for use by this policy. The value must be a multiple of two.

Values 2 to 256

ipv6

Specifies to use **ipv6-criteria**. The software must allocate resources from the chunks allotted to IPv6 criteria.

no-ipv6

Specifies to not use **ipv6-criteria**. Resources are then allocated from the chunk allotted to either IPv4 criteria or MAC criteria, depending on what criteria the user uses.

9.5.1.4 Service ingress QoS policy forwarding class commands

broadcast-meter

Syntax

broadcast-meter *meter-id*
no broadcast-meter

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default broadcast forwarding type meter mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy as a multipoint meter before the mapping can be made. After the FC mapping is executed, all broadcast traffic on a SAP using this policy is forwarded using the *meter-id*.

The broadcast forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of this command reverts the broadcast forwarding type *meter-id* to the default of tracking the multicast forwarding type meter mapping.

Parameters

meter-id

Specifies an existing multipoint queue defined in the **config>qos>sap-ingress** context.

Values 1 to 32

Default 1

meter

Syntax

meter *meter-id*
no meter

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default unicast forwarding type meter mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy as a non-multipoint meter before the mapping can be made. After the FC mapping is executed, all unicast traffic (this includes all traffic, even broadcast and multicast for services) on a SAP using this policy is forwarded using the *meter-id*.

The **no** form of this command reverts the unicast (point-to-point) *meter-id* to the default meter for the FC.

Parameters

meter-id

Specifies an existing non-multipoint meter defined in the **config>qos>sap-ingress** context.

Values 1 to 32

Default 1

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 meter mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy as a multipoint meter before the mapping can be made. After the FC mapping is executed, all multicast traffic on a SAP using this policy is forwarded using the *meter-id*.

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

The **no** form of this command reverts the multicast forwarding type *meter-id* to the default meter for the FC. If the broadcast and unknown forwarding types were not explicitly defined to a multipoint meter, they are also reverted to the default multipoint meter.

Parameters

| | |
|-----------------|--|
| <i>meter-id</i> | Specifies an existing multipoint meter defined in the config>qos>sap-ingress context. |
| Values | 1 to 32 |
| Default | 1 |

unknown-meter

Syntax

unknown-meter *meter-id*
no unknown-meter

Context

config>qos>sap-ingress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides the default unknown unicast forwarding type meter mapping for **fc** *fc-name*. The specified *meter-id* must exist within the policy as a multipoint meter before the mapping can be made. After the FC mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the *meter-id*. The unknown forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of this command reverts the unknown forwarding type *meter-id* to the default of tracking the multicast forwarding type meter mapping.

Parameters

| | |
|-----------------|--|
| <i>meter-id</i> | Specifies an existing multipoint meter defined in the config>qos>sap-ingress context. |
| Values | 1 to 32 |
| Default | 1 |

9.5.1.5 Service ingress QoS policy entry commands

action

Syntax

action [**fc** *fc-name*]

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 FC with specific IP or MAC criteria entry ID. The action command supports setting the FC parameter. Packets that meet all match criteria within the entry have their FC 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 previously entered values for **fc** *fc-name* are overridden with the newly defined parameters.

The **no** form of this command removes the entry from the active entry list. Removing an entry on a policy immediately removes the entry from all SAPs using the policy. All previous parameters for the action are lost.

Default

action specified by the **default-fc** command

Parameters

fc *fc-name*

Specifies the FC name for the queue. When a packet matches the rule, the FC is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the FC is not explicitly defined in the rule, the FC is inherited based on previous rule matches. The value specified for **fc** *fc-name* must be one of the predefined FCs in the system.

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

entry

Syntax

[no] **entry** *entry-id* [**create**]

Context

config>qos>sap-ingress>ip-criteria

config>qos>sap-ingress>mac-criteria

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates or edits an IP or MAC criteria entry for the policy. Multiple entries can be created using unique *entry-id* numbers.

The list of flow criteria is evaluated in a top-down fashion with the lowest entry ID at the top and the highest entry ID at the bottom. If the defined match criteria for an entry within the list matches the information in the egress packet, the system stops matching the packet against the list and performs the matching reclassification actions specified by the entries. If none of the entries match the packet, the IP flow reclassification list has no effect on the packet.

An entry is not populated in the list unless the **action** command is executed for the entry. An entry that is not populated in the list has no effect on egress packets. If the action command is executed without any explicit reclassification actions specified, the entry is populated in the list allowing packets matching the entry to exit the list, preventing them from matching entries lower in the list. Because this is the only flow reclassification entry that the packet matched and this entry explicitly states that no reclassification action is to be performed, the matching packet is not reclassified.

The **no** form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services where that policy is applied.

Parameters

entry-id

Specifies a match criterion and the corresponding action. It is recommended that multiple entries be assigned 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 may not have any match criteria defined (in which case, everything matches) but it must have at least the command **action fc** for it to be considered complete. Entries without the **action** command are considered incomplete and are rendered inactive.

Values 1 to 64

create

This keyword is required when creating a flow entry if 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 configures match criteria for SAP QoS policy matches. When the match criteria are satisfied, the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, then all criteria must be satisfied (and function) before the action associated with the match will be executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

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

Parameters

protocol *protocol-id*

Specifies an IP protocol to be used as an ingress or egress network QoS policy match criterion.

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

Values protocol-id: 0 to 255 protocol numbers accepted in decimal, hexadecimal, or binary

keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-frag, ipv6-icmp, ipv6-no-nxt, isis, ipv6-opts, isoip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

* — udp/tcp wildcard

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

This command creates or edits the match MAC criteria for ingress SAP QoS policy match criteria. When the match criteria are satisfied the action associated with the match criteria is executed.

If more than one match criterion (within one match statement) is 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*.

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 are satisfied, the action associated with the match criteria is executed.

If more than one match criterion (within one match statement) is 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 specified per entry.

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

Parameters

next-header next-header

Specifies the next header to match.

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

Values protocol numbers accepted in DHB: 0 to 42, 45 to 49, 52 to 59, 61 to 255

keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, isoip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

* — udp/tcp wildcard

9.5.1.6 IP QoS policy match commands

dscp

Syntax

dscp *dscp-value* | *dscp-name* [*dscp-mask*]

no dscp

Context

config>qos>sap-ingress>ip-criteria>entry>match

config>qos>sap-ingress>ipv6-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a DiffServ Code Point (DSCP) used for classification of packets from the specified FC.

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



Note:

- This feature is applicable for **ip-criteria** (**any** and **dscp-only**) and **ipv6-criteria** (**any** and **dscp-only**).
- The user is not be allowed to configure dscp name and dscp mask combinations.
- When the user configures dscp value alone, the "show" command displays dscp value as configured value and dscp mask as "FC".
- Use of dscp-value and dscp-mask allows for efficient use of match resources in hardware, because the specification of mask allows user to combine individual DSCP entries to a single value/mask pair and specify similar action for all of them.

Parameters

dscp-value

Specifies the DSCP value in hexadecimal, decimal, or binary format.

Values 0 to 64

dscp-name

Specifies a DSCP name that has been previously mapped to a value using the *dscp-name* command. The DSCP can only be specified by its name.

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

dscp-mask

Specifies a 6-bit mask that can be configured using the formats described in the following table.

This parameter is only supported on the 7210 SAS-D and 7210 SAS-Dxp.

Table 38: DSCP mask value format

| Format style | Format syntax | Example |
|--------------|---------------|---------|
| Decimal | D | 4 |
| Hexadecimal | 0xH | 0x4 |
| Binary | 0bBBB | 0b100 |

To select a range from 4 up to 7 specify 4 and 0b000100 for value and mask.

Values 0 to 64 (decimal, hexadecimal, or binary)

Default 64 (exact match)

dst-ip

Syntax

dst-ip {ip-address/mask}
dst-ip {ipv6-address/mask}
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 use as a SAP QoS policy match criterion.

To match on the destination address, specify the address and its associated mask; for example, 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

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

Parameters

ip-address

Specifies the IP address of the destination IP interface. This address must be unique within the subnet and specified in dotted-decimal notation.

Values ipv4-prefix - a.b.c.d

ipv6-address

Specifies the IP address of the destination IPv6 interface. This address must be unique within the subnet.

Values ipv6-prefix - x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:d.d.d.d
 x - 0 to FFFF (hexadecimal)
 d - 0 to 255 (decimal)

mask

Specifies the length in bits of the subnet mask.

Values 1 to 32 (IPv4)
 1 to 128 (IPv6; 7210 SAS-D)
 1 to 64 (IPv6; 7210 SAS-Dxp)

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

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

Parameters

eq *dst-port-number*

Specifies the destination TCP or UDP port number, specified as equal to (eq) the destination port value specified as a decimal integer.

Values 1 to 65535

ip-prec

Syntax

ip-prec *ip-prec-value* [*ip-prec-mask*]

no ip-prec

Context

config>qos>sap-ingress>ip-criteria>entry>match

config>qos>sap-ingress>ipv6-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies an IP precedence value that must be matched to perform the associated classification actions. If an ingress packet on the SAP where the SAP ingress QoS policy is applied to matches the specified IP precedence value, the actions associated with this entry are taken.

The *ip-prec-value* is derived from the most significant three bits in the IP header ToS byte field (precedence bits). The three precedence bits define eight Class-of-Service (CoS) values commonly used to map packets to per-hop Quality-of-Service (QoS) behavior. The precedence bits are also part of the newer DSCP method of mapping packets to QoS behavior. The DSCP uses the most significant six bits in the IP header ToS byte and so overlaps with the precedence bits.

Both IP precedence and DSCP classification rules are supported. A match entry cannot match on both IP DSCP and IP precedence, which means that the user can use either IP DSCP or IP precedence match in a match entry but not both. The software blocks configuration of ip-precedence match if the **ip-dscp** command is configured already. The converse is also true. A single policy having multiple match entries can have entries such that some of them match IP DSCP and others match IP precedence. The order of the entry determines the priority of the match.

The **no** form of this command removes the IP precedence match criterion.

Parameters

ip-prec-value

Specifies the unique IP header ToS byte precedence bits value that matches the IP precedence rule.

Values 0 to 7

ip-prec-mask

Specifies a mask that can be configured using the formats described in the following table.

Table 39: IP precedence mask value format

| Format style | Format syntax | Example |
|--------------|---------------|---------|
| Decimal | D | 4 |
| Hexadecimal | 0xH | 0x4 |
| Binary | 0bBBB | 0b100 |

To select a range from 4 up to 7, specify an *ip-prec-value* of 4 and an *ip-prec-mask* of 0b100 for value and mask.

src-ip

Syntax

src-ip *ipv4-address/mask*

no src-ip

Context

config>qos>sap-ingress>ip-criteria>entry>match
config>qos>sap-egress>ip-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a source IPv4 address range to use as a SAP QoS policy match criterion.

To match on the source IPv4 address, specify the address and its associated mask; for example, 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

The **no** form of this command removes the source IPv4 address match criterion.

Parameters

ipv4-address

Specifies the IPv4 address of the source IP interface. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

mask

Specifies the subnet mask length, expressed as an integer or in dotted-decimal notation.

Values 1 to 32

src-ip

Syntax

src-ip *ipv6-address/mask*

no src-ip

Context

config>qos>sap-ingress>ipv6-criteria>entry>match
config>qos>sap-egress>ipv6-criteria>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a source IPv6 address range to use as a SAP QoS policy match criterion. The **no** form of this command removes the source IPv6 address match criterion.

Parameters

ipv6-address

Specifies the IPv6 address of the source IP interface. This address must be unique within the subnet.

| | | |
|---------------|--------------|-------------------------------------|
| Values | ipv6-address | x:x:x:x:x:x:x (eight 16-bit pieces) |
| | | x:x:x:x:x:d.d.d.d |
| | | x - 0 to FFFF (hexadecimal) |
| | | d - 0 to 255 (decimal) |

mask

Specifies the subnet mask length, expressed as an integer.

Values 1 to 128 (7210 SAS-D)
1 to 64 (7210 SAS-Dxp)

src-port

Syntax

src-port {eq} *src-port-number*
no src-port

Context

config>qos>sap-ingress>ip-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 SAP QoS policy match criterion.
The **no** form of this command removes the source port match criterion.

Parameters

eq *src-port-number*

Specifies the source TCP or UDP port number, as a decimal integer, against which to match equal (eq) values.

Values 1 to 65535

9.5.1.7 Service ingress MAC QoS policy match commands

dot1p

Syntax

dot1p *dot1p-value* [*dot1p-mask*]
no dot1p

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the IEEE 802.1p value to use as the match criterion.
The **no** form of this command removes the dot1p value as the match criterion.

Parameters

dot1p-value

Specifies the IEEE 802.1p value, expressed as a decimal integer.

Values 0 to 7

dot1p-mask

Specifies a 3-bit mask that can be configured using the following formats.

Table 40: 3-bit mask format

| Format style | Format syntax | Example |
|--------------|---------------|---------|
| Decimal | D | 4 |
| Hexadecimal | 0xH | 0x4 |
| Binary | 0bBBB | 0b100 |

Values 1 to 7 (decimal)

Default 7 (decimal; exact match)

dst-mac

Syntax

dst-mac *ieee-address* [*ieee-address-mask*]
no dst-mac

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a destination MAC address or range to use 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 use 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 formats described in the following table.

Table 41: dot1p mask value format

| 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 0xFFFFFFFF000000

Values 0x0000000000000000 to 0xFFFFFFFFFFFFFFF (hexadecimal)

Default 0xFFFFFFFFFFFFFFF (hexadecimal; exact match)

etype

Syntax

- etype etype-value*
- no etype*

Context

config>qos>sap-ingress>mac-criteria>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an Ethernet type II value for to 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.

For the 7210 SAS-D and 7210 SAS-Dxp, 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 Ethertype value to be used as a match criterion in decimal or hexadecimal.

Values 0x0600 to 0xFFFF

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 that can be configured using the formats in the following table.

Table 42: MAC mask format

| Format style | Format syntax | Example |
|--------------|------------------|-----------------|
| Decimal | DDDDDDDDDDDDDDDD | 281474959933440 |
| Hexadecimal | 0xHHHHHHHHHHHHHH | 0x0FFFFFF000000 |
| Binary | 0bBBBBBBB...B | 0b11110000...B |

To configure all packets with a source MAC OUI value of 00-03-FA are subject to a match condition, then the entry should be specified as: 003FA000000 0xFFFFFFFF000000

Values 0x000000000000000 to 0xFFFFFFFFFFFFFFF (hexadecimal)

Default 0xFFFFFFFFFFFFFFF (hexadecimal; exact match)

9.5.1.8 Service meter QoS policy commands

meter

Syntax

meter *meter-id* [**multipoint**] [**create**]

no meter *meter-id*

Context

config>qos>sap-ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an ingress SAP QoS policy meter.

This command allows the creation of multipoint meters. Only multipoint meters can receive ingress packets that need flooding to multiple destinations. By separating the unicast for multipoint traffic at service ingress and handling the traffic on separate multipoint meters, special handling of the multipoint traffic is possible. Each meter acts as an accounting and (optionally) policing device offering precise control over potentially expensive multicast, broadcast, and unknown unicast traffic. Only the back-end support of multipoint traffic (between the FC and the meter based on forwarding type) needs to be defined. The individual classification rules used to place traffic into FCs are not affected. Meters must be defined as multipoint within the policy at the time of creation.

The multipoint meters are for service traffic destined for multiple destinations, such as multicast traffic in a VPLS service. Within non-multipoint services, such as Epipe services, all traffic is considered unicast because of the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service is mapped to a multipoint service meter.

When an ingress SAP QoS policy with multipoint meters is applied to an Epipe SAP, the multipoint meters are not created.

Any billing or statistical queries about a multipoint meter on a non-multipoint service return zero values. Any meter parameter information requested about a multipoint meter on a non-multipoint service returns the meter parameters in the policy. Multipoint meters are not created for non-multipoint services.

The **no** form of this command removes the *meter-id* from the SAP ingress QoS policy and from any existing SAPs using the policy. Any FC mapped to the meter reverts to the default meters. When a meter is removed, any pending accounting information for each SAP meter, created as a result of the definition of the meter in the policy, is discarded.

Parameters

meter-id

Specifies the meter within the policy. This is a required parameter each time the **meter** command is executed.

Values 1 to 32

multipoint

Keyword that creates the meter as a multipoint meter.

create

Keyword to create a meter.

adaptation-rule

Syntax

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

no adaptation-rule

Context

config>qos>sap-ingress>meter

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 settings when the meter is provisioned in hardware. For the CIR and PIR parameters, individually the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of 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 **pir** and **cir** apply.

Default

adaptation-rule cir closest pir closest

Parameters

cir adaptation-rule

Specifies the adaptation rule used while computing the operational CIR value and defines the constraints enforced when adapting 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 for the meter. When the **cir** parameter is not specified, the default constraint applies.



Note:

For 7210 SAS-Dxp, see [Adaptation rule for meters](#) for information about calculating the next multiple equal to or less than the specified rate.

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. When **min** is defined, the operational CIR is the next multiple of 8 kbps that is greater than or equal to the specified rate.

closest — Specifies that the operational CIR value is equal to the closest specified rate, taking into account the hardware step size. When **closest** is defined, the operational CIR is the next multiple of 8 kbps that is closest to the specified rate.

pir adaptation-rule

Specifies the adaptation rule used while computing the operational PIR value and defines the constraints enforced when adapting 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 for the meter. When the rate command is not specified, the default applies.



Note:

For 7210 SAS-Dxp, see [Adaptation rule for meters](#) for information about calculating the next multiple equal to or less than the specified rate.

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. When **min** is defined, the operational PIR is the next multiple of 8 kbps that is greater than or equal to the specified rate.

closest — Specifies that the operational PIR value is equal to the closest specified rate, taking into account the hardware step size. When **closest** is defined, the operational PIR is the next multiple of 8 kbps that is closest to the specified rate.

cbs

Syntax

cbs *size* [kbits | bytes | kbytes]
no cbs

Context

config>qos>sap-ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command provides a mechanism to override the default CBS for the meter. The **cbs** command specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the CBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter configured parameters.



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

The **no** form of this command returns the CBS size to the default value.

Default

32 kbits

Parameters

size

Specifies the number of kilobits, kilobytes, or bytes reserved for the meter, expressed as an integer. For example, if a value of 100 kb is wanted, enter the value 100 and the keyword kbits. The bucket size is rounded off to the next highest 4096-byte boundary.

Values kbits — 4 to 16384, default (7210 SAS-D)
 4 to 2146959, default (7210 SAS-Dxp)
 bytes — 512 to 2097152, default (7210 SAS-D)
 512 to 274810752, default (7210 SAS-Dxp)
 kbytes — 1 to 2048, default (7210 SAS-D)
 1 to 268369, default (7210 SAS-Dxp)

color-mode

Syntax

color-mode *color-mode*

no color-mode

Context

config>qos>sap-ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the meter to operate in either color-aware or color-blind mode.

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

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

The following behavior is available in color-aware mode:

- If the packet is pre-colored as in-profile (or green packets), depending on the burst size of the packet, the meter can either be marked in-profile or out-profile.

- If the packet is pre-colored as out-profile (or yellow packets), even if the packet burst is less than the current available CBS, it is not marked as in-profile and remains as out-profile.
- If the packet burst is higher than the MBS, it is marked as red and is dropped by the meter at ingress.

Default

color-blind

Parameters

color-mode

Specifies the mode the meter operates in.

Values color-aware — The meter operates in color-aware mode.
 color-blind —The meter operates in color-blind mode.

mbs

Syntax

mbs *size* [kbits | bytes | kbytes]

no mbs

Context

config>qos>sap-ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a mechanism to override the default MBS for the meter. The MBS parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the MBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying with configured meter parameters.



Note:

The adaptation rule configured for the rate influences the step-size used for the burst. See [Adaptation rule for meters](#) for information.

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

Default

128 kbits

Parameters

size

Specifies the size reserved for the meter. For example, if a value of 100 kb is wanted, enter the value 100 and the keyword kbits. The bucket size is rounded off to the next highest 4096-byte boundary.

Values kbits — 4 to 16384, default (7210 SAS-D)
4 to 2146959, default (7210 SAS-Dxp)
bytes — 512 to 2097152, default (7210 SAS-D)
512 to 274810752, default (7210 SAS-Dxp)
kbytes — 1 to 2048, default (7210 SAS-D)
1 to 268369, default (7210 SAS-Dxp)

mode

Syntax

mode {trtcm1 | trtcm2 | srtcm}

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 (trTCM1) or Single Rate Three Color Marker (srTCM). The mode command can be executed at anytime.



Note:

- The meter counters are reset to zero when the meter mode is changed.
- For more information about the interpretation of rate parameters when the meter mode is configured as "trtcm2", see the command description of the **policer rate** command.

The **no** form of the command sets the default mode **trtcm1**.

Default

trtcm1

Parameters

trtcm1

Specifies the policing algorithm defined in RFC2698 and meters the packet stream and marks its packets as green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked as yellow or green, depending on whether it exceeds the CIR. The trTCM1 is useful for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate. Two token buckets are used, the CBS bucket and the MBS bucket. Tokens are added to the buckets based on the CIR and PIR rates. The algorithm deducts tokens from both the CBS and the MBS buckets to determine a profile for the packet.

trtcm2

Specifies the policing algorithm defined in RFC4115 and meter the packet stream and marks its packets either green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked yellow or green, depending on whether it exceeds the CIR. The trTCM2 is useful for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate. Two token buckets are used, the CBS bucket and the EBS bucket. Tokens are added to the buckets based on the CIR and EIR rates. The algorithm deducts tokens from either the CBS bucket (that is, when the algorithm identifies the packet as in-profile or green packet) or the EBS bucket (that is, when the algorithm identifies the packet as out-of-profile or yellow packet).



Note:

When the meter mode is configured in trtcm2 mode, the system interprets the PIR rate parameter as EIR and the MBS parameter as the EBS value for use by RFC 4115 algorithm.

srtcm

Specifies that an IP packet IP packet stream will be metered and marks its packets green, yellow, or red. Marking is based on a CIR and two associated burst sizes, a CBS and an MBS. A packet is marked green if it does not exceed the CBS, yellow if it exceeds the CBS but not the MBS, and red otherwise. The srTCM is useful for ingress policing of a service where only the length, not the peak rate, of the burst determines service eligibility.

rate

Syntax

rate *cir* *cir-rate-in-kbps* [**pir** *pir-rate-in-kbps*]

no rate

Context

config>qos>sap-ingress>meter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the administrative PIR and CIR parameters for the meter.

The rate command can be executed at anytime, altering the PIR and CIR rates for all meters created through the association of the SAP Ingress QoS policy with the *meter-id*.

The **no** form of this command reverts all meters created with the *meter-id* by association with the QoS policy to the default PIR(max) and CIR(0) parameters.

Default

rate cir 0 pir max — The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the *pir-rate* value.

Parameters

cir cir-rate-in-kbps

Specifies that the **cir** parameter overrides the default administrative CIR used by the meter. When the rate command has not been executed, or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed, and must be specified as a positive integer.

The actual CIR rate is dependent on the meter's **adaptation-rule** parameters and the hardware.

Values 0 to 20000000, max (7210 SAS-Dxp)
0 to 4000000, max (7210 SAS-D)

pir pir-rate-in-kbps

Specifies the administrative PIR rate, in kilobits, for the meter. When this command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed. When the **rate** command is executed, a PIR setting is optional. Fractional values are not allowed, and must be specified as a positive integer.

The actual PIR rate is dependent on the meter's **adaptation-rule** parameters and the hardware.

Values 0 to 20000000, max (7210 SAS-Dxp)
0 to 4000000, max (7210 SAS-D)

9.5.1.9 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.

Values 1 to 65535

Default all SAP ingress policies

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 outputs are examples of QoS SAP ingress policy information, and [Table 43: Output fields: QoS SAP ingress](#) describes the output fields:

- [Sample output](#)
- [Sample output for 7210 SAS-D with DSCP value and mask](#)
- [Sample output for 7210 SAS-D with DSCP name](#)

Sample output

```
*A:SAS>show>qos# sap-ingress 1 detail

=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (1)
-----
Policy-id           : 1           Scope           : Template
Default FC          : be
Criteria-type       : None
Accounting           : packet-based
Classifiers Allowed  : 4           Meters Allowed   : 2
Classifiers Reqrđ (VPLS) : 2       Meters Reqrđ (VPLS) : 2
Classifiers Reqrđ (EPIPE) : 1     Meters Reqrđ (EPIPE) : 1
Description          : Default SAP ingress QoS policy.
```

```
-----
Meter Mode   CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS Admin  MBS Admin
          CIR Oper                PIR Oper  CBS Oper  MBS Oper
-----
1      TrTcm1    0         closest    max        closest    7 KBits  10 KBits
          0         max
11     TrTcm1    0         closest    max        closest    8 KBits  11 KBits
          0         max
          def      def

-----
FC              UCastM      MCastM      BCastM      UnknownM
-----
No FC-Map Entries Found.

-----
Match Criteria
-----
No Matching Criteria.

-----
Associations
-----
No Associations Found.

=====
*A:SAS>show>qos#
```

The following output is an example of SAP ingress QoS policy information about the 7210 SAS-D when the DCSP value and mask are configured.

Sample output for 7210 SAS-D with DSCP value and mask

```
*A:7210 SAS> show qos sap-ingress 2 detail

=====
QoS Sap Ingress
=====

Sap Ingress Policy (2)
-----
Policy-id           : 2                Scope           : Template
Default FC          : be
Criteria-type       : IP                IP-Mac Rule Priority : None
Mac Sub-Criteria    : None              IP Sub-Criteria     : any
IPv6 Enabled        : False             IPv6 Sub-Criteria    : dscp
Accounting          : packet-based
Classifiers Allowed : 2                Meters Allowed      : 1
Classifiers Reqr'd (VPLS) : 2          Meters Reqr'd (VPLS) : 1
Classifiers Reqr'd (L3 Mc) : 2          Meters Reqr'd (L3 Mc) : 1
Classifiers Reqr'd (EPIPE) : 2          Meters Reqr'd (EPIPE) : 1
Description         : (Not Specified)

-----
Meter Mode   CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS Admin  MBS Admin
          Color Mode  CIR Oper                PIR Oper  CBS Oper  MBS Oper
-----
1      TrTcm1    0         closest    max        closest    def      def
          color-blind 0         200000000  32      512

-----
FC              UCastM      MCastM      BCastM      UnknownM
-----
```

```

No FC-Map Entries Found.

-----
Match Criteria
-----
IP Match Criteria
-----
Entry                : 1
Description          : (Not Specified)
Source IP            : Undefined          Source Port          : None
Dest. IP             : Undefined          Dest. Port            : None
Protocol             : none              DSCP value/mask       : 4/5
Fragment             : Off               Ip Precedence         : None
FC                   : be

-----
IPv6 Match Criteria
-----
No Match Criteria Entries found.

SAP Associations
-----
Service-Id           : 1 (Epipe)          Customer-Id           : 1
- SAP : 1/1/24

```

Sample output for 7210 SAS-D with DSCP name

The following output is an example of SAP ingress QoS policy information about the 7210 SAS-D when the DSCP name is configured.

```

*A:>config>service>epipe>sap$ /show qos sap-ingress 2 detail

=====
QoS Sap Ingress
=====
Sap Ingress Policy (2)
-----
Policy-id             : 2                  Scope                : Template
Default FC            : be
Criteria-type         : IP                IP-Mac Rule Priority  : None
Mac Sub-Criteria      : None              IP Sub-Criteria       : any
IPv6 Enabled          : False             IPv6 Sub-Criteria     : dscp
Accounting            : packet-based
Classifiers Allowed   : 2                  Meters Allowed       : 1
Classifiers Reqrđ (VPLS) : 2              Meters Reqrđ (VPLS)  : 1
Classifiers Reqrđ (L3 Mc) : 2             Meters Reqrđ (L3 Mc) : 1
Classifiers Reqrđ (EPIPE) : 2             Meters Reqrđ (EPIPE) : 1
Description           : (Not Specified)

-----
Meter Mode    CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS Admin  MBS Admin
Color Mode    CIR Oper   PIR Oper   CBS Oper   MBS Oper
-----
1    TrTcm1    0          closest   max        closest   def        def
    color-blind 0          200000000
-----
FC      UCastM    MCastM    BCastM    UnknownM
-----
No FC-Map Entries Found.

```

```

-----
Match Criteria
-----
IP Match Criteria
-----
Entry                : 1
Description          : (Not Specified)
Source IP            : Undefined          Source Port        : None
Dest. IP             : Undefined          Dest. Port          : None
Protocol             : none              DSCP                : cp63
Fragment             : Off               Ip Precedence       : None
FC                   : be
-----
IPv6 Match Criteria
-----
No Match Criteria Entries found.

SAP Associations
-----
Service-Id           : 1 (Epipe)          Customer-Id          : 1
- SAP : 1/1/24

```

Table 43: Output fields: QoS 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 |
| Default FC | Specifies the default FC 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 FC for matched traffic |
| | MAC — Specifies that a MAC criteria-based SAP is used to select the appropriate ingress meters and corresponding FC for matched traffic |
| 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 the 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 |

| Label | Description |
|----------------------|---|
| Classifiers Allowed | Indicates the number of classifiers allowed for a service |
| Classifiers Required | Indicates the number of classifiers for a VPLS or Epipe service |
| Description | A text string that helps identify the policy's context in the configuration file |
| Meter | Displays the meter ID |
| Mode | Specifies the configured mode of the meter (trTcm1, trTcm2 or srTcm) |
| CIR Admin | Specifies the administrative 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 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 |

| Label | Description |
|----------------------|---|
| 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 |
| Entry | The entry ID in a policy or filter table |
| DSCP | Specifies a DiffServ Code Point (DSCP) name used for an ingress SAP QoS policy match |
| FC | Specifies the entry's FC |
| Src MAC | Specifies a source MAC address or range to be used as a Service Ingress QoS policy match |
| Dst MAC | Specifies a destination MAC address or range to be used as a Service Ingress QoS policy match |
| dot1p | Specifies a IEEE 802.1p value to be used as the match |
| Ethernet-type | Specifies an Ethernet type II Ethertype value to be used as a Service Ingress QoS policy match |
| FC | Specifies the entry's FC |
| Service Association | |
| Service-Id | The unique service ID number which identifies the service in the service domain |
| Customer-Id | Specifies the customer ID which identifies the customer to the service |
| SAP | Specifies the a Service Access Point (SAP) within the service where the SAP ingress policy is applied |
| 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 |

10 Access egress QoS policies

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

10.1 Overview

An access-egress policy defines the queuing for the traffic egressing on the access ports. Access-egress queue policies are used at the Ethernet access port and define the bandwidth distribution for the various FC/queue traffic egressing on the Ethernet access port.

There is one default access egress policy which is identified as policy ID 1. Each policy has eight queues available. The FC-to-queue mapping is predefined and cannot be changed. Queue parameters, for example CIR, PIR, and so on, can be modified. The default policy can be copied but cannot be deleted or modified.

10.1.1 Basic configurations

A basic access egress QoS policy must conform to the following:

- have a unique access egress QoS policy ID
- have a QoS policy scope of **template** or **exclusive**
- have queue parameters that can be modified but not deleted

10.1.1.1 Modifying access egress QoS queues

To modify access-egress queue parameters, specify the following:

- **queue ID value**

Eight queues are identified and are mapped as defined in [Table 15: FC-to-queue ID map](#).

- **queue parameters**

Egress queues support configuration of CIR and PIR rates.

Example

The following is a sample access egress QoS policy configuration output.

```
A:card-1>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
....
    access-egress 30 create
        remarking
        queue 1
            rate cir 100 pir 4500
        exit
```

```
        queue 2
        exit
        queue 3
        exit
        queue 4
        exit
        queue 5
        exit
        queue 6
        exit
        queue 7
        exit
        queue 8
        exit
    exit
-----
A:card-1>config>qos#
```

10.1.1.2 Applying access egress QoS policies

Apply access egress policies to the 'Ethernet ports' entity.

A policy can be applied to the ports that are in access mode.

10.1.1.2.1 Ethernet ports

Use the following CLI syntax to apply an access-egress policy to an Ethernet port.

```
config>port#
    ethernet access egress
        qos access-egress-policy-id
```

Example

The following is a sample port configuration output.

```
*A:card-1>config>port# info
-----
        shutdown
        ethernet
            access
                egress
                    qos 30
            exit
        exit
    exit
-----
*A:card-1>config>port#
```


10.1.1.3 Default access egress QoS policy values

Example

The following are sample access egress default policy parameters.

```
*A:card-1>config>qos>access-egress# info detail
-----
description "Default Access egress QoS policy."
no remarking
scope template
queue 1
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 2
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 3
    adaptation-rule cir closest pir closest
    rate 0 pir max
exit
queue 4
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 5
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 6
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 7
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
queue 8
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
exit
-----
*A:card-1>config>qos>access-egress#
```

The following table lists the default FC marking values when remarking is enabled on the access egress policy.

Table 44: Default FC marking values for 7210 SAS-D and 7210 SAS-Dxp

| Default FC value | dot1p marking values |
|------------------|---|
| af: | dot1p-in-profile 2 dot1p-out-profile 2 |
| be: | dot1p-in-profile 0 dot1p-out-profile 0 |

| Default FC value | dot1p marking values |
|------------------|---|
| ef: | dot1p-in-profile 5 dot1p-out-profile 5 |
| h1: | dot1p-in-profile 6 dot1p-out-profile 6 |
| h2: | dot1p-in-profile 4 dot1p-out-profile 4 |
| l1: | dot1p-in-profile 3 dot1p-out-profile 3 |
| l2: | dot1p-in-profile 1 dot1p-out-profile 1 |
| nc: | dot1p-in-profile 7 dot1p-out-profile 7 |

10.1.1.4 Deleting QoS policies

Every access Ethernet port is associated, by default, with the default access-egress policy (*policy-id* 1). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the port configuration. When you remove a non-default access-egress policy, the association reverts to the default *policy-id* 1.

A QoS policy cannot be deleted until it is removed from all access ports where it is applied.

When a QoS policy is still in use, the following error occurs.

```
*A:card-1>config>qos# no access-egress 30
MINOR: CLI Could not remove Access egress policy "30" because it is in use.
```

10.1.1.5 Removing a policy from the QoS configuration

Use the following syntax to remove a policy.

```
config>qos# no access-egress policy-id
```

Example

```
config>qos# no access-egress 100
config>qos# no access-egress 1010
```

10.2 Access egress QoS policy command reference

10.2.1 Command hierarchies

- [Configuration commands for 7210 SAS-D and 7210 SAS-Dxp](#)
- [Show commands](#)

10.2.1.1 Configuration commands for 7210 SAS-D and 7210 SAS-Dxp

```
- config
- qos
  - access-egress policy-id [create]
  - no access-egress policy-id
    - description description-string
    - no description
    - fc fc-name [create]
    - no fc fc-name
      - [no] de-mark [force de-value]
      - dot1p dot1p-priority
      - no dot1p
      - dot1p-in-profile dot1p-value
      - no dot1p-in-profile
      - dot1p-out-profile dot1p-value
      - no dot1p-out-profile
      - dscp-in-profile dscp-name
      - no dscp-in-profile
      - dscp-out-profile dscp-name
      - no dscp-out-profile
    - queue queue-id
      - adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
      - no adaptation-rule
      - rate cir cir-rate [pir pir-rate]
      - no rate
    - [no] remarking {use-dot1p | use-dscp | all}
    - scope {exclusive | template}
    - no scope
```

10.2.1.2 Show commands

```
- show
- qos
  - access-egress [policy-id] [association| detail]
```

10.3 Command descriptions

10.3.1 Configuration commands

10.3.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>access-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.

access-egress

Syntax

access-egress *policy-id* [**create**]

no access-egress *policy-id*

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates or edits an access-egress QoS policy. The policy can be applied to multiple access ports. The access-egress policy is common to services (SAPs) all configured on that access port.

Any changes made to an existing policy are applied to all access ports on which the policy is specified.

The remarking parameters and queue parameters are used when port-based queuing is configured.

Parameters

policy-id

Specifies the value that uniquely identifies the access-egress policy.

Values 1 to 65535

create

Specifies that an access-egress policy is created. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

fc

Syntax

fc *fc-name* [**create**]

no fc *fc-name*

Context

config>qos>access-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the **fc** node within the access-egress QoS policy is used to contain the explicitly defined dot1p marking commands for the *fc-name*.

If the mapping for the *fc-name* and marking value is not defined, the node for *fc-name* is not displayed in the show configuration or save configuration output.

The **no** form of the command removes the explicit dot1p marking commands for the *fc-name*.

Parameters

fc-name

Specifies this FC for which dot1p marking is to be edited. The value specified for *fc-name* must be one of the predefined FCs in the system.

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

create

Specifies that an access-egress policy is created. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

de-mark

Syntax

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

Context

```
config>qos>access-egress>fc  
config>qos>network>egress>fc
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

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

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

Parameters

de-value

Specifies the DEI value to use for this FC.

Values 0 or 1

dot1p

Syntax

[no] dot1p *dot1p-value*

Context

```
config>qos>access-egress>fc  
config>qos>network>egress>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.

DEI marking can be enabled using the de-mark command along with this command for the command to take effect. When de-mark command is configured along with this command, then the DEI bit is marked in the packet to indicate the profile of the packet. The DEI bit is marked to 0 to indicate in-profile/green packet and 1 to indicate out-of-profile/yellow packet. If the 'force de-value' parameter is specified then the DEI bit is set to specified value for all packets.

If the **no** form of this command is executed the software will use the dot1p-in-profile and dot1p-out-profile if configured, else it will use default values.



Note:

The following rules are applied by the software to determine the dot1p values to when both **dot1p** command and **dot1p-in-profile** and dot1p-out-profile command is specified.

1. If de-mark is not configured, the dot1p [in|out]-profile values are considered. Even if **dot1p dot1p-value** is configured, it is ignored. If the dot1p [in|out]-profile value is not configured, the default values are considered for that FC.
2. If de-mark is configured and if **dot1p dot1p-value** is configured then it is considered. Else if the dot1p [in|out]-profile value is configured it is considered. In this case **dot1p dot1p-value** has the precedence over dot1p [in|out]-profile.

Default

no dot1p

Parameters

dot1p-value

Specifies the 802.1p value to set for in-profile frames in this FC.

Values 0 to 7

dot1p-in-profile

Syntax

```
dot1p-in-profile dot1p-priority  
no dot1p-in-profile
```

Context

config>qos>access-egress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

The command enables the context to mark on an egress the in and out of 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/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with in-profile status (or green color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the dot1p command has no effect.

If DEI marking is enabled using the de-mark command and the command **dot1p dot1p-value** is used to configure the dot1p value, then this command has no effect. In other words, enabling DEI marking has precedence over this command and the system ignores this command.

When this command is used the DEI Bit is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DEI bit is set to 0.

The **no** form of this command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.



Note:

The following rules are applied by the software to determine the dot1p values to when both the **dot1p** command and the **dot1p-in-profile** and **dot1p-out-profile** command is specified:

- If de-mark is not configured, the dot1p [in|out]-profile values are considered. Even if **dot1p dot1p-value** is configured it is ignored. If the dot1p [in|out]-profile value is not configured, the default values are considered for that FC.
- If de-mark is configured, and if the **dot1p dot1p-value** command is configured, then it is considered. Else if 'dot1p [in|out]-profile' value is configured it is considered. In this case **dot1p dot1p-value** has the precedence over dot1p [in|out]-profile.
- If marking is enabled and both **dot1p** and dot1-[in|out]-profile commands are not specified, then the value reverts to the default.

Default

0

Parameters

dot1p-priority

Specifies the 802.1p value to set for in-profile frames in this FC.

Values 0 to 7

dot1p-out-profile

Syntax

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

Context

config>qos>access-egress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

The command enables the context to mark on an egress the in and out of 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/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets with out-of-profile status (or yellow color) of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the dot1p command has no effect.

If DEI marking is enabled using the de-mark command and the dot1p-value is configured, then this command has no effect. In other words, enabling DEI marking has precedence over this command and the system ignores this command.

When this command is used the DEI Bit is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DEI bit is set to 0.

The **no** form of this command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.



Note:

The following rules are applied by software to determine the dot1p values to when both dot1p command and **dot1p-in-profile** and **dot1p-out-profile** command is specified:

- If de-mark is not configured, the dot1p [in|out]-profile values are considered. Even if **dot1p dot1p-value** command is configured, it is ignored. If the dot1p [in|out]-profile value is not configured, the default values are considered for that FC.
- If de-mark is configured and if **dot1p dot1p-value** command is configured, then it is considered. Else if 'dot1p [in|out]-profile' value is configured, it is considered. In this case **dot1p dot1p-value**, has the precedence over dot1p [in|out]-profile.
- If marking is enabled and both **dot1p dot1p-value** and dot1p-[in|out]-profile commands are not specified, the value revert to the default.

Default

0

Parameters

dot1p-priority

Specifies the 802.1p value to set for in-profile frames in this FC.

Values 0 to 7

dscp-out-profile

Syntax

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

Context

config>qos>access-egress>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 FC. 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 FC queue.

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

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

Parameters

dscp-name

Specifies the DSCP name.

Values be | cp1 | cp2 | cp3 | cp 4| cp5 | cp6 | cp7| cs1 | cp9 | af11 | cp11 | af12
| cp13 | af13 | cp15 | cs2 | cp17 | af21 | cp19 | af22 | cp21 | af23 | cp23
| cs3 | cp25 | af31 | cp27 | af32 | cp29 | af33 | cp31 | cs4 | cp33 | af41 |
cp35 | af42 | cp37 | af43 | cp39 | cs5 | cp41 | cp42 | cp43 | cp44 | cp45
| ef | cp47 | nc1 | cp49 | cp50 | cp51 | cp52 | cp53 | cp54 | cp55 | nc2 |
cp57 | cp58 | cp59 | cp60 | cp61 | cp62 | cp63

dscp-in-profile

Syntax

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

Context

config>qos>access-egress>fc

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the in-profile DSCP name for the FC. 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 FC queue.

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

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

Parameters

dscp-name

Specifies the DSCP name.

Values

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

queue

Syntax

queue *queue-id*

Context

config>qos>access-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

Parameters

queue-id

Specifies the access-egress queue-id associated with an FC according to [Table 15: FC-to-queue ID map](#).

Values 1 to 8

remarking

Syntax

[no] remarking {use-dot1p | use-dscp | all}

Context

config>qos>access-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command remarks egress packets sent out of access ports and access-uplink ports. For 7210 SAS-D, remarking can be enabled or disabled.



Note:

If remarking is enabled without specifying one of the options, by default 'use-dot1p' is used for access-egress and "all" is used for network-egress.

The **no** form of this command disables remarking, which is the default behavior.

Default

no remarking

Parameters

use-dot1p

Specifies that the dot1p bits are marked in the packet header for all IEEE 802.1q and IEEE 802.1p encapsulated traffic sent out of the access port.

use-dscp

Specifies that the IP DSCP bits are marked in the packet header for IPv4 traffic sent out of the access port.

all

Specifies that the dot1p bits are marked in the packet header for all IEEE 802.1q and IEEE 802.1p encapsulated traffic, and in addition the IP DSCP bits are marked in the packet header for all IPv4 traffic sent out the access port.

10.3.1.2 Access egress queue QoS policy commands

adaptation-rule

Syntax

adaptation-rule [*cir adaptation-rule*] [*pir adaptation-rule*]
no adaptation-rule

Context

config>qos>access-egress>queue

Platforms

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 settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of 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.

rate

Syntax

rate cir cir-rate [pir pir-rate]

no rate

Context

config>qos>access-egress>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the administrative PIR and 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 any time, altering the PIR and CIR rates for all queues created on the access ports.

The **no** form of this command reverts all queues created using the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (max, 0).

Parameters

cir-rate

Specifies that the default administrative CIR used by the queue is overridden. When the **rate** command is executed, a valid CIR setting must be explicitly defined. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be specified as a positive integer.

Values For devices with only 1G ports: 0 to 1000000, max
For devices with both 1G and 10G ports: 0 to 10000000, max

Default 0

pir-rate

Specifies the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a PIR setting is optional. When the **rate** command has not been executed, the default PIR of **max** is assumed. Fractional values are not allowed and must be specified as a positive integer.

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

Values For devices with only 1G ports: 0 to 1000000, max
For devices with both 1G and 10G ports: 0 to 10000000, max

Default max

scope

Syntax

scope {exclusive | template}

no scope

Context

config>qos>access-egress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the scope as exclusive or template. The policy's scope cannot be changed if the policy is applied to multiple ports.

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

Default

template

Parameters

exclusive

Specifies that the policy can only be applied to one 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 becomes available for assignment to another exclusive interface. The system default policies cannot be put into the **exclusive** scope. An error occurs if scope exclusive is executed in the default access-egress policy (*policy-id* 1).

template

Specifies that the policy can be applied to multiple ports on the router.

access-egress

Syntax

access-egress [*policy-id*] [**association** | **detail**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays access-egress QoS policy information.

Parameters

policy-id

Displays the policy ID of the access-egress policy.

association

Displays associations related to the specified access-egress policy.

detail

Displays detailed policy information including the policy associations.

Output

The following output is an example of access-egress QoS policy information, and [Table 45: Access-egress labels and descriptions](#) describes the output fields.

Sample output

```
A:Dut-A>show>qos# access-egress

=====
Access Egress Policies
=====
Policy-Id      Scope      Description
-----
1              Template  Default Access egress QoS policy.
```



```

=====
A:Dut-A>show>qos#
A:Dut-A>show>qos# access-egress 1 detail
=====
QoS Access Egress
=====
-----
Policy-id      : 1                               Scope      : Template
Remark        : False                           Remark Pol Id: 2
Description    : Default Access egress QoS policy.
-----
Queue Rates and Rules
-----
-----
QueueId      CIR      CIR Adpt Rule      PIR      PIR Adpt Rule
-----
Queue1       0        closest           max      closest
Queue2       0        closest           max      closest
Queue3       0        closest           max      closest
Queue4       0        closest           max      closest
Queue5       0        closest           max      closest
Queue6       0        closest           max      closest
Queue7       0        closest           max      closest
Queue8       0        closest           max      closest
-----
Queue Mode and Weight Details
-----
-----
QueueId      Mode      Weight
-----
Queue1       weighted  1
Queue2       weighted  1
Queue3       weighted  1
Queue4       weighted  1
Queue5       weighted  1
Queue6       weighted  1
Queue7       weighted  1
Queue8       weighted  1
-----
High Slope
-----
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down       70            90          75
Queue2       Down       70            90          75
Queue3       Down       70            90          75
Queue4       Down       70            90          75
Queue5       Down       70            90          75
Queue6       Down       70            90          75
Queue7       Down       70            90          75
Queue8       Down       70            90          75
-----
Low Slope
-----
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down       50            75          75
Queue2       Down       50            75          75
Queue3       Down       50            75          75
Queue4       Down       50            75          75
Queue5       Down       50            75          75

```

| | | | | |
|-------------------------------------|------|-----|---------------------|------------|
| Queue6 | Down | 50 | 75 | 75 |
| Queue7 | Down | 50 | 75 | 75 |
| Queue8 | Down | 50 | 75 | 75 |
| ----- | | | | |
| Burst Sizes and Time Average Factor | | | | |
| ----- | | | | |
| QueueId | CBS | MBS | Time Average Factor | Queue-Mgmt |
| ----- | | | | |
| Queue1 | def | def | 7 | default |
| Queue2 | def | def | 7 | default |
| Queue3 | def | def | 7 | default |
| Queue4 | def | def | 7 | default |
| Queue5 | def | def | 7 | default |
| Queue6 | def | def | 7 | default |
| Queue7 | def | def | 7 | default |
| Queue8 | def | def | 7 | default |
| ----- | | | | |
| Associations | | | | |
| ----- | | | | |
| ===== | | | | |
| A:Dut-A>show>qos# | | | | |
| A:Dut-A>show>qos# | | | | |

Table 45: Access-egress labels and descriptions

| Label | Description |
|-----------------------|--|
| Scope | Exclusive — Specifies that this policy can be applied only to a single access-egress port |
| | Template — Specifies that this policy can be applied to multiple access ports on the router |
| Description | A text string that helps identify the policy context in the configuration file |
| Queue Rates and Rules | |
| QueueId | Displays the queue identifier associated with the SAP-egress QoS policy |
| CIR | Specifies the administrative CIR parameters for the queue. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth |
| CIR Adpt Rule | min — The operational CIR for the queue is equal to or greater than the administrative rate specified using the rate command |
| | max — The operational CIR for the queue is less than or equal to the administrative rate specified using the rate command |
| | closest — The operational CIR for the queue is the rate closest to the rate specified using the rate command without exceeding the operational PIR |
| PIR | Specifies the administrative PIR parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the access port |

| Label | Description |
|-------------------------------------|--|
| PIR Adpt Rule | min — The operational PIR for the queue is equal to or greater than the administrative rate specified using the rate command |
| | max — The operational PIR for the queue is less than or equal to the administrative rate specified using the rate command |
| | closest — The operational PIR for the queue is the rate closest to the rate specified using the rate command |
| High Slope/Low slope | |
| QueueId | Displays the queue identifier associated with the SAP-egress QoS policy |
| State | Displays the state of the queue The state of the queue can be either "Up" or "Down" |
| Start Avg | Specifies the low priority or high priority RED slope position for the shared buffer average utilization value, where the packet discard probability starts to increase above zero |
| Max Avg | Specifies the percentage of the shared buffer space for the buffer pool, at which point the drop probability becomes "1" This parameter is expressed as a decimal integer |
| Max Prob | Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one |
| Burst Sizes and Time Average Factor | |
| QueueId | Displays the queue identifier associated with the SAP-egress QoS policy |
| CBS | Displays the configured CBS value |
| MBS | Displays the configured MBS value |
| Time Average Factor | Displays the value of the time average factor in use |

11 Port scheduler policies

This chapter provides information to configure port scheduler policies using the command line interface.

11.1 Configuring port scheduler policies

The **port-scheduler-policy** command creates a port scheduler template which may be assigned to an egress port. Only one port scheduler policy is allowed per port. There is a "default" port-scheduler policy (which services the queues of the port in a Strict order) associated with each port. To change the behavior, users can associate the port with another port-scheduler policy. The policy contains mode commands to set the mode of scheduling (RR, Strict, WRR, WDRR) and queue commands to set the weight of the queue (only 8 queues per port and queue settings only for WRR/WDRR modes). In WRR/WDRR, a **strict** option treats that particular queue as a strict queue, this leads to a hybrid mode of scheduling (WRR+Strict, WDRR+Strict).

11.2 Basic configurations

A basic QoS port scheduler policy must conform to the rule that each QoS port scheduler policy must have a unique policy name.

11.2.1 Creating a QoS port scheduler policy

To create a port scheduler policy, define the following:

- a port scheduler policy name
- include a description - the description provides a brief overview of policy features

Use the following syntax to create a QoS port scheduler policy.

Note that the **create** keyword is included in the command syntax upon creation of a policy.

A port scheduler policy cannot be deleted unless it is removed from all ports where it is applied. The "default" port-scheduler policy cannot be deleted.

```
config>qos
  port-scheduler-policy port-scheduler-name [create]
    description description-string
    mode {strict | rr | wrr | wdr}
    queue queue-id [strict | weight weight]
```

Example

The following is a sample port scheduler policy configuration output.

```
*A:card-1>config>qos>port-sched-plcy# info
-----
mode WRR
```

```

        queue 1 weight 1
        queue 2 weight 3
        queue 3 weight 5
        queue 5 weight 5
        queue 6 weight 1
-----
*A:card-1>config>qos>port-sched-plcy#

```

11.3 Service management tasks

This section describes the service management tasks.

11.3.1 Copying and overwriting port scheduler policies

You can copy an existing QoS policy, rename it with a new QoS policy value, or overwrite an existing policy. The overwrite option must be specified or an error occurs if the destination policy exists.

```
config>qos> copy port-scheduler-policy src-name dst-name [overwrite]
```

Example

```

config>qos# copy port-scheduler-policy psp psp1

*A:card-1>config# qos port-scheduler-policy psp create
*A:card-1>config>qos>port-sched-plcy# mode WRR
*A:card-1>config>qos>port-sched-plcy# queue 1 weight 1
*A:card-1>config>qos>port-sched-plcy# queue 2 weight 3
*A:card-1>config>qos>port-sched-plcy# queue 3 weight 5
*A:card-1>config>qos>port-sched-plcy# exit
*A:card-1>config# qos copy port-scheduler-policy psp psp1
*A:card-1>config# qos copy port-scheduler-policy psp psp1
MINOR: CLI Destination "psp1" exists - use {overwrite}.
*A:card-1>config# qos copy port-scheduler-policy psp psp1 overwrite
*A:card-1>config# show qos port-scheduler-policy
=====
Port Scheduler Policies
=====
Policy-Id      Description                                     Mode
-----
default        Default Port Scheduler policy.                  STRICT
psp            WRR
psp1           WRR
=====
*A:card-1>config#

*A:card-1>config# show qos port-scheduler-policy psp
=====
QoS Port Scheduler Policy
=====
Policy-Name : psp
Accounting: packet-based
Mode : WRR
Last changed : 01/01/2000 22:13:01
Queue 1 Weight: 1
Queue 2 Weight: 3
Queue 3 Weight: 5

```

```
Queue 4 Weight: strict
Queue 5 Weight: strict
Queue 6 Weight: strict
Queue 7 Weight: strict
Queue 8 Weight: strict
=====
*A:card-1>config#
*A:card-1>config# show qos port-scheduler-policy psp1
=====
QoS Port Scheduler Policy
=====
Policy-Name : psp1
Accounting: packet-based
Mode : WRR
Last changed : 01/01/2000 22:13:17
Queue 1 Weight: 1
Queue 2 Weight: 3
Queue 3 Weight: 5
Queue 4 Weight: strict
Queue 5 Weight: strict
Queue 6 Weight: strict
Queue 7 Weight: strict
Queue 8 Weight: strict
=====
*A:card-1>config#
```

11.3.2 Editing QoS policies

You can edit a port-scheduler policy, the modifications are done and it affects the port where it is applied. The "default" port-scheduler policy cannot be modified.

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.

11.4 QoS port scheduler policy command reference

11.4.1 Command hierarchies

- [Port scheduler policy configuration commands](#)
- [Operational commands](#)
- [Show commands](#)

11.4.1.1 Port scheduler policy configuration commands

```
- config
- qos
- [no] port-scheduler-policy port-scheduler-name [create]
  - description description-string
  - no description
  - mode {strict | rr | wrr | wdrr}
  - no mode
```

```
- queue queue-id [strict | weight weight]  
- no queue queue-id
```

11.4.1.2 Operational commands

```
- config  
  - qos  
    - copy port-scheduler-policy src-name dst-name [overwrite]
```

11.4.1.3 Show commands

```
- show  
  - qos  
    - port-scheduler-policy [port-scheduler-policy-name] [association]
```

11.5 Command descriptions

11.5.1 Configuration commands

11.5.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>port-scheduler-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.

11.5.1.2 Operational commands

copy

Syntax

copy port-scheduler-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 port scheduler QoS policy entries for a port scheduler QoS policy to another port scheduler QoS policy. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters

port-scheduler-policy src-name dst-name

Specifies that the source policy and the destination policy are port scheduler policy IDs.
Specifies the source policy that the copy command will attempt to copy from and specify the destination policy name to which the command will copy a duplicate of the policy.

overwrite

Specifies that everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, a message is generated saying that the destination policy ID exists

11.5.1.3 Port scheduler policy commands

port-scheduler-policy

Syntax

[**no**] port-scheduler-policy *port-scheduler-name* [**create**]

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the port-scheduler policy. The default scheduling done for a port is strict scheduling. When a port-scheduler policy is applied to a port, it overrides the default scheduling and determines the type of scheduling (Strict, RR, WRR, WDRR, WRR/WDRR + Strict) to be done between the 8 CoS queues of that particular port. When a port scheduler policy is detached from a port, the port reverts back to the default scheduling (strict).

The **no** form of this command removes the policy from the system.

Parameters

port-scheduler-name

Specifies an existing policy name. Each port-scheduler policy name should be unique and can go up to 32 ASCII characters.

create

Specifies that a port scheduler policy is created.

mode

Syntax

mode {strict | rr | wrr | wdrr}

no mode

Context

config>qos>port-sched-plcy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a particular mode of scheduling for the policy. For example, this implies that when a policy with a mode RR is applied to a port then that port will follow the round robin type of scheduling between its queues.

Parameters

strict

Specifies that the port follows a strict scheduler mode.

rr

Specifies that the port follows round robin scheduling.

wrr

Specifies that the port follows weighted round robin scheduling.

wdr

Specifies that the port follows weighted deficit round robin scheduling.

queue

Syntax

queue *queue-id* [**strict** [**weight** *weight*]

no queue *queue-id*

Context

config>qos>port-sched-plcy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a port scheduler queue. The queue and its weights can be configured only for WRR/WDRR modes. The weight specified in case of WRR corresponds to the number of packets that needs to be sent out in a cycle for that particular queue.

For WDRR, the weight specified is the ratio of traffic that will be sent out for that particular queue. For example, in WDRR, if a weight value for queue 1 is 1 and a weight value for queue 2 is 5, then traffic out of the port is in the ratio of 1:5 between the queues (1 and 2) provided no traffic is flowing in the other queues. If the keyword **strict** is specified in any of the queues, then that particular queue will be treated as strict. This set of strict priority queues is serviced first in the order of their CoS numbering (the higher numbered CoS queue receives service before smaller numbered queues).

The **no** form of this command under a WRR/WDRR mode will set the queue weights to default (for example, 1).

Parameters

queue-id

Specifies the queue ID.

Values 1 to 8 (8 is the highest)

strict

Specifies strict access.

weight *weight*

Specifies the number of packets in case of WRR and ratio of traffic out in WDRR.

Values 1 to 15

11.5.1.4 Show commands

port-scheduler-policy

Syntax

port-scheduler-policy [*port-scheduler-policy-name*] [**association**]

Context

show>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays port-scheduler policy information

Parameters

- port-scheduler-policy-name**
Displays information for the specified existing port scheduler policy.
- association**
Displays associations related to the specified port scheduler policy.

Output

The following outputs are examples of port scheduler policy information, and [Table 46: Output fields: show port scheduler policy](#) describes the output fields.

Sample output

```
*A:card-1>config# show qos port-scheduler-policy
=====
Port Scheduler Policies
=====
Policy-Id      Description                                     Mode
default       Default Port Scheduler Policy.                STRICT
psp           psp                                           WRR
test          psp                                           WRR
=====
*A:card-1>config#
*A:card-1>config# show qos port-scheduler-policy psp association
=====
QoS Port Scheduler Policy
=====
Policy-Name : psp
Mode : WRR
Accounting: packet-based
-----
Associations
-----
- Port : 1/1/1
```

```
=====
*A:card-1>config#
*A:card-1>config# show qos port-scheduler-policy psp
=====
QoS Port Scheduler Policy
=====
Policy-Name : psp
Mode : WRR
Accounting : packet-based
Last changed : 01/01/2000 05:14:06
Queue 1:      Weight: 1
Queue 2:      Weight: 3
Queue 3:      Weight: 5
Queue 4:      Weight: 0
Queue 5:      Weight: 5
Queue 6:      Weight: 5
Queue 7:      Weight: strict
Queue 8:      Weight: strict
=====
*A:card-1>config#
*A:SN12345678>config# show qos port-scheduler-policy default
=====
QoS Port Scheduler Policy
=====
Policy-Name      : default
Accounting       : frame-based
Description      : Default Port Scheduler policy.
Mode            : STRICT
Last changed     : 08/04/2009 20:55:46
Number Of Queues : 8
=====
*A:SN12345678>config#
```

Sample output for 7210 SAS-D

```
*A:SAS-D>show>qos# port-scheduler-policy abc

=====
QoS Port Scheduler Policy
=====
Policy-Name      : abc
Description      : (Not Specified)
Accounting       : packet-based
Mode            : STRICT
Last changed     : 01/01/1970 04:57:48
Number Of Queues : 8
=====
*A:SAS-D>show>qos# port-scheduler-policy abc association

=====
QoS Port Scheduler Policy
=====
Policy-Name      : abc
Description      : (Not Specified)
Accounting       : packet-based
Mode            : STRICT

-----
Associations
-----
No Association Found.
```

```
=====
*A: SAS-D>show>qos#
```

Table 46: Output fields: show port scheduler policy

| Label | Description |
|--------------|--|
| Policy Name | Displays the port scheduler policy name |
| Associations | Displays associations related to the specified port scheduler policy |
| Mode | Displays the port scheduler policy mode (STRICT, RR, WRR, WDRR) |
| Accounting | Displays whether the accounting mode is frame-based or packet-based |
| Last Changed | Displays the last time the configuration changed |
| Queue # | Displays the weight of the queue if configured |

12 Slope QoS policies

This chapter provides information to configure slope QoS policies using the command line interface.

12.1 Overview of buffer pools

Default buffer pool exists (logically) at each port. Buffer pools cannot be created or deleted in the 7210 SAS. The egress buffer pools are created as access uplink port egress buffer pool for access-uplink ports and access port egress buffer pool for access ports. Based on the maximum number of ports to be supported for access and access-uplink, the total buffer is distributed into the access port egress buffer pool and the access uplink port egress buffer pool. The distribution of the buffers into access and access-uplink port egress pools take care of the buffer requirements at the port level for various queue shaping/scheduling mechanisms and for various packet sizes varying from 64 bytes to jumbo frames. Each port on the system gets an equal portion of the available buffers. From the buffers allocated to a port, each queue gets its CBS amount of buffers. The remaining buffers are allocated toward the shared MBS pool per port. All the queues of the port can use the buffers from the shared MBS pool. By default, each queue on the access port and access-uplink port is associated with slope-policy default which disables the slope parameters within the pool.

On 7210 SAS-D, WRED is supported to evaluate the packet's eligibility to be allocated a buffer based on the slope parameters configured for the queue.

On 7210 SAS-Dxp, WRED is supported to evaluate the TCP packet's eligibility to be allocated a buffer based on the slope parameters configured for the queue. WRED is not supported for non-TCP packets.

12.1.1 Configuration guidelines for 7210 SAS-D

7210 SAS-D provides an option to use either 2 slope per queue or 3 slopes per queue. This can be configured using the CLI command **config>system>qos>no use-wred-slopes**. The option to use only 2 WRED slopes per queue (port egress queues), allows differentiating in-profile and out-of-profile traffic flows. The option to use 3 WRED slopes per queues allows differentiating in-profile TCP traffic, out-of-profile TCP traffic and non-TCP traffic (both in and out-profile use a single slope). The slope does not get enabled by default. To maintain backward compatibility, by default the system uses 3 slopes option and user has to change it explicitly to use 2 slopes (if they need).

The following table compares the WRED slope used for different traffic flows with 2 slopes or 3 slopes per queue.

Table 47: Slope behavior (applicable to 7210 SAS-D)

| Slopes | TCP-non-TCP slope option (uses 3 WRED slopes per queue) | High-low slope option (uses 2 WRED slopes per queue) |
|---|--|--|
| SAP Ingress TCP/IP traffic (Number of VLAN tags <=2) | High-priority TCP slope or low-priority TCP slope, based on packet profile | High-priority or low-priority slope, based on packet profile |

| Slopes | TCP-non-TCP slope option (uses 3 WRED slopes per queue) | High-low slope option (uses 2 WRED slopes per queue) |
|--|--|--|
| SAP Ingress non-TCP traffic (Number of VLAN tags does not matter) | Non-TCP slope - No in/out profile differentiation | High-priority or low-priority slope, based on packet profile |
| SAP Ingress TCP/IP traffic (Number of VLAN tags>2) | Non-TCP slope - No in/out profile differentiation | High-priority or low-priority slope, based on packet profile |

12.1.2 Configuration guidelines for 7210 SAS-Dxp

7210 SAS-Dxp can use two slopes per queue. Using two WRED slopes per queue (port egress queues) allows differentiation of in-profile and out-of-profile TCP traffic flows. The slope is not enabled by default.



Note:

All non-TCP traffic is tail-dropped if the egress queues are full when the traffic is being enqueued.

The following table describes the WRED slope used for different traffic flows with two slopes per queue.

Table 48: Slope behavior (applicable to 7210 SAS-Dxp)

| Slopes | High-low slope option (Uses 2 WRED slopes per queue) |
|--|--|
| SAP Ingress TCP/IP traffic (Number of VLAN tags <=2) | High-priority or low-priority slope, based on packet profile |
| SAP Ingress non-TCP traffic (Number of VLAN tags does not matter) | Tail-drop |
| SAP Ingress TCP/IP traffic (Number of VLAN tags>2) | Tail-drop |

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, low slope and non-TCP slope are shut down (default).
- Default values can be modified but parameters cannot be deleted.

12.2.1 Create a slope QoS policy for 7210 SAS-D

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 for 7210 SAS-D devices, define the following:

- **a slope policy ID value**

The system will not dynamically assign a value.

- **include a description**

The description provides a brief overview of policy features.

- Option to use either 2 slopes per queue or 3 slopes per queue.
- 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.
- Optional non-TCP slope for the non-TCP Random Early Detection (RED) slope graph.
- The time average factor (TAF), a weighting exponent used to determine the portion of the shared buffer instantaneous utilization and shared buffer average utilization used to calculate the new shared buffer average utilization.

Use the following CLI syntax to configure a slope policy for both 2-slope per queue and 3-slope per queue option. When using 2-slope per queue, the non-tcp slope parameters are not used.

```
config>qos
  slope-policy name
    description description-string
    high-slope
      start-avg percent
      max-avg percent
      max-prob percent
      no shutdown
    low-slope
      start-avg percent
      max-avg percent
      max-prob percent
      no shutdown
    non-tcp-slope
      start-avg percent
      max-avg percent
      max-prob percent
      no shutdown
    time-average-factor taf
```

The following displays the slope policy configuration for 7210 SAS-D:

```
A:ALA-7>config>qo>slope-policy# info
-----
      description "slope policy SlopePolicy1"
      high-slope
        no shutdown
      exit
      low-slope
        no shutdown
      exit
      non-tcp-slope
        no shutdown
```



```
exit
-----
A:ALA-7>config>qos>slope-policy#
```

12.2.2 Create a slope QoS policy for 7210 SAS-Dxp

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 for 7210 SAS-Dxp devices, define the following:

- a slope policy ID value; the system will not dynamically assign a value
- a description; the description provides a brief overview of policy features
- two slopes per queue:
 - the high slope for the high priority Random Early Detection (RED) slope graph
 - the low slope for the low priority Random Early Detection (RED) slope graph
- the time average factor (TAF), a weighting exponent used to determine the portion of the shared buffer instantaneous utilization and shared buffer average utilization used to calculate the new shared buffer average utilization

Use the following CLI syntax to configure a slope policy with two slopes per queue.

```
config>qos
  slope-policy name
    description description-string
    high-slope
      start-avg percent
      max-avg percent
      max-prob percent
      no shutdown
    low-slope
      start-avg percent
      max-avg percent
      max-prob percent
      no shutdown
    time-average-factor taf
```

The following displays the slope policy configuration for 7210 SAS-Dxp:

```
A:ALA-7>config>qo>slope-policy# info
-----
      description "slope policy SlopePolicy1"
      high-slope
        no shutdown
      exit
      low-slope
        no shutdown
      exit
      exit
-----
A:ALA-7>config>qos>slope-policy#
```

12.3 Applying slope policies

Based on the 7210 SAS platform capabilities, the slope policies are associated with different entities:

Example

The following CLI syntax examples may be used to apply slope policies to access ports or access-uplink ports.

```
config>port>access>egress>pool>slope-policy name
config>port>access>uplink>egress>pool>slope-policy name
```

12.3.1 Default slope policy values

This section provides the default values for slope policies.

12.3.1.1 Default slope values for 7210 SAS-D

The default access egress and access uplink egress policies are identified as *policy-id* "default". The default policies cannot be edited or deleted. The following table lists default policy parameters.

Table 49: Slope policy defaults for 7210 SAS-D

| Field | Default |
|----------------------|----------------------|
| description | Default slope policy |
| high (RED) slope | |
| Administrative state | shutdown |
| start-avg | 70% utilization |
| max-avg | 90% utilization |
| max-prob | 75% |
| low (RED) slope | |
| Administrative state | shutdown |
| start-avg | 50% utilization |
| max-avg | 75% utilization |
| max-prob | 75% |
| non-TCP (RED) slope | |
| Administrative state | shutdown |

| Field | Default |
|-----------|-----------------|
| start-avg | 50% utilization |
| max-avg | 75% utilization |
| max-prob | 75% |

```

A:ALA>config>qos# slope-policy default
A:ALA>config>qos>slope-policy# info detail
-----
description "Default slope policy."
queue "1"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit
  non-tcp-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit
  time-average-factor 7
exit
queue "2"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit
  non-tcp-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit
  time-average-factor 7
exit
queue "3"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope

```

```
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    non-tcp-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    time-average-factor 7
exit
queue "4"
    high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
    exit
    low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    non-tcp-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    time-average-factor 7
exit
queue "5"
    high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
    exit
    low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    non-tcp-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    time-average-factor 7
exit
queue "6"
    high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
    exit
    low-slope
        shutdown
```

```

        start-avg 50
        max-avg 75
        max-prob 75
    exit
    non-tcp-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit
queue "7"
    high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    non-tcp-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit
queue "8"
    high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    non-tcp-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit

```

```
-----
A:ALA>config>qos>slope-policy#
```

12.3.1.2 Default slope values for 7210 SAS-Dxp

The default access egress and access uplink egress policies are identified as *policy-id* "default". The default policies cannot be edited or deleted. The following table lists default policy parameters.

Table 50: Slope policy defaults for 7210 SAS-Dxp

| Field | Default |
|----------------------|----------------------|
| description | Default slope policy |
| high (RED) slope | |
| Administrative state | shutdown |
| start-avg | 70% utilization |
| max-avg | 90% utilization |
| max-prob | 75% |
| low (RED) slope | |
| Administrative state | shutdown |
| start-avg | 50% utilization |
| max-avg | 75% utilization |
| max-prob | 75% |

```

A:ALA>config>qos# slope-policy default
A:ALA>config>qos>slope-policy# info detail
-----
description "Default slope policy."
queue "1"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit
time-average-factor 7
exit
queue "2"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit

```

```
        time-average-factor 7
    exit
    queue "3"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "4"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "5"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "6"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "7"
        high-slope
```

```
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
    exit
    low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    time-average-factor 7
exit
queue "8"
    high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
    exit
    low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    time-average-factor 7
exit
```

A:ALA>config>qos>slope-policy#

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* default. A QoS policy cannot be deleted until it is removed from all ports where it is applied or if the policies are using the slope-policy.

```
ALA-7>config>qos# no slope-policy slopePolicy1
MINOR: QOS #1902 Slope policy has references
ALA-7>config>qos#
```

Example

The following CLI syntax examples can be used to remove slope policies from ports.

```
config>port>access>egress>pool# no slope-policy name
config>port>access>uplink>egress>pool# no slope-policy name
```


12.4.1 Remove a policy from the QoS configuration

To delete a slope policy, enter the following command:

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

Example

```
config>qos# no slope-policy slopePolicy1
```

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]
```

The following output displays the copied policies for 7210 SAS-D devices.

```
A:ALA-7210>config>qos#
-----
...
    description "Default slope policy."
      queue "1"
        high-slope
          shutdown
          start-avg 70
          max-avg 90
          max-prob 75
        exit
        low-slope
          shutdown
          start-avg 50
          max-avg 75
          max-prob 75
        exit
        non-tcp-slope
          shutdown
          start-avg 50
          max-avg 75
          max-prob 75
        exit
        time-average-factor 7
      exit
    queue "2"
      high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
      exit
      low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
```

```
        exit
        non-tcp-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
        exit
        time-average-factor 7
    exit
    queue "3"
        high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
        exit
        low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
        exit
        non-tcp-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
        exit
        time-average-factor 7
    exit
    queue "4"
        high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
        exit
        low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
        exit
        non-tcp-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
        exit
        time-average-factor 7
    exit
    queue "5"
        high-slope
        shutdown
        start-avg 70
        max-avg 90
        max-prob 75
        exit
        low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
```

```
        non-tcp-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    time-average-factor 7
exit
queue "6"
    high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    non-tcp-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit
queue "7"
    high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    non-tcp-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    time-average-factor 7
exit
queue "8"
    high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
    exit
    low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
    exit
    non-tcp-slope
```

```
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    time-average-factor 7
exit
...
-----
A:ALA-7210>config>qos#
```

The following output displays the copied policies for 7210 SAS-Dxp devices.

```
A:ALA-7210>config>qos#
-----
...
    description "Default slope policy."
    queue "1"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "2"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
    queue "3"
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 75
        exit
        time-average-factor 7
    exit
exit
```

```
queue "4"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit
  time-average-factor 7
exit
queue "5"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit
  time-average-factor 7
exit
queue "6"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit
  time-average-factor 7
exit
queue "7"
  high-slope
    shutdown
    start-avg 70
    max-avg 90
    max-prob 75
  exit
  low-slope
    shutdown
    start-avg 50
    max-avg 75
    max-prob 75
  exit
  time-average-factor 7
exit
queue "8"
  high-slope
    shutdown
    start-avg 70
```

```
        max-avg 90
        max-prob 75
    exit
    low-slope
        shutdown
        start-avg 50
        max-avg 75
        max-prob 75
    exit
    time-average-factor 7
exit
...
-----
A:ALA-7210>config>qos#
```

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

12.7.1 Command hierarchies

- [Configuration commands for 7210 SAS-D](#)
- [Configuration commands for 7210 SAS-Dxp](#)
- [WRED commands \(supported only on 7210 SAS-D\)](#)
- [Operational commands](#)
- [Show commands](#)

12.7.1.1 Configuration commands for 7210 SAS-D

```
- config
- qos
- slope-policy name [create]
- no slope-policy name
- description description-string
- no description
- queue queue-id
- [no] high-slope
- max-avg percent
- no max-avg
- max-prob percent
- no max-prob
- [no] shutdown
- start-avg percent
- no start-avg
- [no] low-slope
```

```
- max-avg percent
- no max-avg
- max-prob percent
- no max-prob
- [no] shutdown
- start-avg percent
- no start-avg
- [no] shutdown
- [no] non-tcp-slope
- max-avg percent
- no max-avg
- max-prob percent
- no max-prob
- [no] shutdown
- start-avg percent
- no start-avg
- time-average-factor value
- no time-average-factor
```

12.7.1.2 Configuration commands for 7210 SAS-Dxp

```
- config
- qos
- slope-policy name [create]
- no slope-policy name
- description description-string
- no description
- queue queue-id
- [no] high-slope
- max-avg percent
- no max-avg
- max-prob percent
- no max-prob
- [no] shutdown
- start-avg percent
- no start-avg
- [no] low-slope
- max-avg percent
- no max-avg
- max-prob percent
- no max-prob
- [no] shutdown
- start-avg percent
- no start-avg
- [no] shutdown
- time-average-factor value
- no time-average-factor
```

12.7.1.3 WRED commands (supported only on 7210 SAS-D)

```
- config
- system
- qos
- no use-wred-slopes
- use-wred-slopes slope-type
```

12.7.1.4 Operational commands

```
- config
  - qos
    - copy slope-policy src-name dst-name [overwrite]
```

12.7.1.5 Show commands

```
- show
  - qos
    - slope-policy [slope-policy-name] [detail]
```

12.8 Command descriptions

12.8.1 Configuration commands

12.8.1.1 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.

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. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters

slope-policy

Specifies that the source policy ID and the destination policy ID are slope policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

overwrite

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

12.8.2 Operational commands

12.8.2.1 Slope policy QoS commands

slope-policy

Syntax

[no] **slope-policy** *name* [**create**]

Context

config>qos

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This configures a QoS slope policy.

Default

slope-policy "default"

Parameters

name

Specifies the name of the slope policy. Valid names consist of any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

12.8.2.2 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 RED slope graph. Each buffer pool supports a high priority RED slope for managing access to the shared portion of the buffer pool for high priority or in-profile packets.

The **high-slope** parameters can be changed at any time and the affected buffer pool high priority RED slopes will be adjusted appropriately.



Note:

See [Table 47: Slope behavior \(applicable to 7210 SAS-D\)](#) for information about the mapping of traffic types to use high-slope parameters on 7210 SAS-D. See [Table 48: Slope behavior \(applicable to 7210 SAS-Dxp\)](#) for information about the mapping of traffic types to use high-slope parameters on 7210 SAS-Dxp.

The **no** form of this command reverts the high slope configuration commands to the default values. If the commands within **high-slope** are set to the default parameters, the **high-slope** node will not appear in save config and show config output unless the detail parameter is present.

low-slope

Syntax

[no] **low-slope**

Context

config>qos>slope-policy>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the low priority RED slope graph. Each buffer pool supports a low priority RED slope for managing access to the shared portion of the buffer pool for low priority or out-of-profile packets.

The **low-slope** parameters can be changed at any time and the affected buffer pool low priority RED slopes must be adjusted appropriately.



Note:

See [Table 47: Slope behavior \(applicable to 7210 SAS-D\)](#) for information about the mapping of traffic types to use low-slope parameters on 7210 SAS-D. See [Table 48: Slope behavior \(applicable to 7210 SAS-Dxp\)](#) for information about the mapping of traffic types to use low-slope parameters on 7210 SAS-Dxp.

The **no** form of this command reverts the low slope configuration commands to the default values. If the leaf commands within **low-slope** are set to the default parameters, the **low-slope** node will not appear in save config and show config output unless the detail parameter is present.

non-tcp-slope

Syntax

[no] non-tcp-slope

Context

config>qos>slope-policy>queue

Platforms

7210 SAS-D

Description

This command configures non-tcp profile RED slope parameters.



Note:

See [Table 47: Slope behavior \(applicable to 7210 SAS-D\)](#) for information about the mapping of traffic types to use non-TCP-slope parameters on 7210 SAS-D.

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

time-average-factor

Syntax

time-average-factor *value*

no time-average-factor

Context

config>qos>slope-policy>queue

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a weighting factor to calculate the new shared buffer average utilization after assigning buffers for a packet entering a queue. To derive the new shared buffer average utilization, the buffer pool takes a portion of the previous shared buffer average and adds it to the inverse portion of the instantaneous shared buffer utilization. The **time-average-factor** command sets the weighting factor between the old shared buffer average utilization and the current shared buffer instantaneous utilization when calculating the new shared buffer average utilization.

The TAF value applies to all high priority, low priority and non-tcp packets WRED slopes for egress access and network buffer pools controlled by the slope policy.

The **no** form of this command reverts the default setting.

Default

7

Parameters

value

Specifies the TAF, expressed as a decimal integer. The value specified for TAF affects the speed at which the shared buffer average utilization tracks the instantaneous shared buffer utilization. A low value weights the new shared buffer average utilization calculation more to the shared buffer instantaneous utilization, zero using it exclusively. A high value weights the new shared buffer average utilization calculation more to the previous shared buffer average utilization value.

Values 0 to 15

12.8.2.3 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>non-tcp-slope (supported only on 7210 SAS-D devices)

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the low priority or high priority or non-tcp Weighted Random Early Detection (WRED) slope position for the reserved and shared buffer average utilization value where the packet discard probability rises directly to one. The percent parameter is expressed as a percentage of the shared buffer size.

The **no** form of this command reverts the value 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 startavg. If the entered value is smaller than the current value of startavg, an error will occur and no change will take place.

Values 0 to 100

max-prob

Syntax

max-prob *percent*

no max-prob

Context

config>qos>slope-policy>queue>high-slope

config>qos>slope-policy>queue>low-slope

config>qos>slope-policy>queue>non-tcp-slope (supported only on 7210 SAS-D devices)

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command sets 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. 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 setting.

Default

max-prob 75

Parameters

percent

Specifies the maximum drop probability percentage corresponding to the max-avg, expressed as a decimal integer.

Values 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 25, 50, 75, 100

shutdown

Syntax

[no] shutdown

Context

config>qos>slope-policy>high-slope

config>qos>slope-policy>low-slope

config>qos>slope-policy>queue

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command enables or disables the administrative status of the RED 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>non-tcp-slope (supported only on 7210 SAS-D)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets 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. 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.

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 will occur and no change will take place.

Values 0 to 100

12.8.2.4 WRED command for 7210 SAS-D

use-wred-slopes

Syntax

use-wred-slopes *slope-type*

no use-wred-slopes

Context

config>system>qos

Platforms

7210 SAS-D

Description

This command provides the user with the option to use 2 WRED slopes per queue or use 3 WRED slopes per queue. It is a global option which affects all the queues in the system. In other words, user can choose to use either 2 WRED slopes for all queues in the system or 3 WRED slopes for all queues in the system.

Using 3 WRED slopes per queue allows differentiating tcp in-profile traffic, tcp out-of-profile traffic, and non-tcp traffic. For non-tcp traffic both in and out profile use the same slope.

Using 2 WRED slopes per queue allows differentiating in-profile and out-of-profile traffic, without further differentiation of tcp and non-tcp traffic. All traffic, irrespective of tcp or non-tcp traffic, uses either in-profile slope or out-of-profile slope, depending on the profile assigned to the traffic by the ingress meters.



Note:

See [Table 47: Slope behavior \(applicable to 7210 SAS-D\)](#) for information about the mapping of traffic types to use high-slope, low-slope, and non-TCP-slope parameters on 7210 SAS-D.

The **no** form of this command enables the use of 3 WRED slopes per queue.

Default

use-wred-slopes tcp-non-tcp to maintain backward compatibility.

Parameters

slope-type

Specifies the slope policy type.

Values High and Low slope type: When high-low is set, two slopes are used per queue. High priority/In-profile slope for all packets that are classified as in-profile by the ingress meter and Low priority/out-of-profile slope for all packets that are classified as out-of-profile by the ingress meter. The high-priority/in-profile WRED slope uses the values configured under **config>qos>slope-policy>high-slope**. The low-priority/out-of-profile WRED slope uses the values configured under **config>qos>slope-policy>low-slope**. The values configured under non-TCP WRED slope is ignored by the system.

TCP and Non-TXP slope type: The non-TCP WRED slope is used for all packets classified as non-TCP packets on ingress, irrespective of the packet's profile or priority. Packets classified as TCP and determined to be high-priority/in-profile by the ingress meter, uses the high priority TCP WRED slope. This slope uses the values configured under **config>qos>slope-policy>high-slope**. Packets classified as TCP and determined to be low-priority/out-of-profile by the ingress meter, uses the low-priority TCP WRED slope. The low-priority/out-of-profile TCP WRED slope uses the values configured under **config>qos>slope-policy>low-slope**. The non-TCP WRED slope uses the values configured under **config>qos>slope-policy>non-tcp-slope**.

12.8.3 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 outputs are examples of QoS slope policy information, and the associated tables describe the output fields:

- [Sample output](#) , [Table 51: Output fields: slope policy](#)
- [Sample output: detailed \(7210 SAS-D\)](#), [Table 52: Output fields: slope policy detail](#)

Sample output

*A:>config# show qos slope-policy 1

| | | |
|------------------|-------------------------|------------------------|
| ===== | | |
| QoS Slope Policy | | |
| ===== | | |
| Policy | : 1 | |
| ----- | | |
| Utilization | State | Start-Threshold |
| ----- | | |
| High-Slope | Down | 75% |
| Low-Slope | Down | 50% |
| ----- | | |
| Queue | High Slope Drop Rate(%) | Low Slope Drop Rate(%) |
| ----- | | |
| Queue 1 | 6.250000 | 100.000000 |
| Queue 2 | 6.250000 | 100.000000 |
| Queue 3 | 6.250000 | 100.000000 |
| Queue 4 | 6.250000 | 100.000000 |
| Queue 5 | 6.250000 | 100.000000 |

```

Queue 6          6.250000          100.000000
Queue 7          6.250000          100.000000
Queue 8          6.250000          100.000000
=====
*A:>config#
*A:>config# show qos slope-policy 1 detail
=====
QoS Slope Policy
=====
Policy           : 1
-----
Utilization      State      Start-Threshold
-----
High-Slope      Down       75%
Low-Slope       Down       50%
-----
Queue           High Slope Drop Rate(%)    Low Slope Drop Rate(%)
-----
Queue 1          6.250000          100.000000
Queue 2          6.250000          100.000000
Queue 3          6.250000          100.000000
Queue 4          6.250000          100.000000
Queue 5          6.250000          100.000000
Queue 6          6.250000          100.000000
Queue 7          6.250000          100.000000
Queue 8          6.250000          100.000000
-----
Associations
-----
Object Type Object Id   Application      Pool
-----
Port         1/1/1       Acc-Egr         default
=====
*A:>config#

```

Table 51: Output fields: 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 |

| Label | Description |
|-----------|--|
| | Specifies the low priority or high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one. |
| Max Prob. | Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one. |

Sample output: detailed (7210 SAS-D)

```
*A:SAS-D>show>qos# slope-policy abc detail
=====
QoS Slope Policy
=====
Policy      : abc
Description : (Not Specified)
-----
High Slope
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down        70            90           75
Queue2       Down        70            90           75
Queue3       Down        70            90           75
Queue4       Down        70            90           75
Queue5       Down        70            90           75
Queue6       Down        70            90           75
Queue7       Down        70            90           75
Queue8       Down        70            90           75
-----
Low Slope
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down        50            75           75
Queue2       Down        50            75           75
Queue3       Down        50            75           75
Queue4       Down        50            75           75
Queue5       Down        50            75           75
Queue6       Down        50            75           75
Queue7       Down        50            75           75
Queue8       Down        50            75           75
-----
Non Tcp Slope
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down        50            75           75
Queue2       Down        50            75           75
Queue3       Down        50            75           75
Queue4       Down        50            75           75
Queue5       Down        50            75           75
Queue6       Down        50            75           75
Queue7       Down        50            75           75
Queue8       Down        50            75           75
-----
Time Avg Factor
-----
```

```

-----
Queue Id    Time Avg Factor
-----
Queue1      7
Queue2      7
Queue3      7
Queue4      7
Queue5      7
Queue6      7
Queue7      7
Queue8      7

-----
Associations
-----
Object Type Object Id      Application      Pool
-----
No Matching Entries

*A:SAS-D>show>qos#

*A:SAH01-051>show>qos# slope-policy 32 detail

=====
QoS Slope Policy
=====
*A:SAH01-051>show>qos#
*A:SAH01-051>show>qos# slope-policy 32 detail

=====
QoS Slope Policy
=====

```

Table 52: Output fields: slope policy detail

| 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 |

| Label | Description |
|-----------|--|
| | 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 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.

13.1 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

13.2 Ethernet

IEEE 802.1AB, Station and Media Access Control Connectivity Discovery is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

TIA-1057, LLDP for Media endpoint devices is supported on Dxp, Sx/S-1/10GE, and Sx/S-1/10GE-VC

IEEE 802.1ad, Provider Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ag, Connectivity Fault Management is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ah, Provider Backbone Bridges is supported on T(N)

IEEE 802.1ax, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1D, MAC Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1p, Traffic Class Expediting is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1Q, Virtual LANs is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1s, Multiple Spanning Trees is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1w, Rapid Reconfiguration of Spanning Tree is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1X, Port Based Network Access Control is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ab, 1000BASE-T is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ac, VLAN Tag is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ad, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ae, 10 Gb/s Ethernet is supported on Dxp, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ah, Ethernet in the First Mile is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ba, 40 Gb/s and 100 Gb/s Ethernet is supported on R6, R12, and Sx-10/100GE

IEEE 802.3i, Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3u, Fast Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3z, Gigabit Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3af, Power Over Ethernet (PoE) is supported on Dxp, T-ETR, Mxp-ETR, and Sx/S-1/10GE



Note:

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3at, Power Over Ethernet (PoE+) is supported on Dxp, T-ETR, Mxp-ETR, and Sx/S-1/10GE



Note:

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3bt, Power Over Ethernet (PoE++/HPoE) is supported on Dxp



Note:

Only on Dxp-16p and Dxp-24p.

ITU-T G.8032, Ethernet Ring Protection Switching is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ITU-T Y.1731, OAM functions and mechanisms for Ethernet based networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

13.3 EVPN

draft-snr-bess-evpn-proxy-arp-nd-00, Proxy-ARP/ND function in EVPN networks is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:

Sx/S-1/10GE standalone mode only.

RFC 7432, BGP MPLS-Based Ethernet VPN is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:

Sx/S-1/10GE standalone mode only.

draft-ietf-bess-evpn-vpws-14, Virtual Private Wire Service support in Ethernet VPN is supported on Mxp

13.4 Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, Operational management of Loop Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

With Segment Routing.

RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:
With Segment Routing.

RFC 7490, Remote Loop-Free Alternate (LFA) Fast Reroute (FRR) is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:
With Segment Routing.

13.5 Internet Protocol (IP) — General

draft-grant-tacacs-02, The TACACS+ Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-vrrp-unified-spec-02, Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6 is supported on Mxp

RFC 768, User Datagram Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 793, Transmission Control Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 854, TELNET Protocol Specifications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 951, Bootstrap Protocol (BOOTP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1034, Domain Names - Concepts and Facilities is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1035, Domain Names - Implementation and Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1350, The TFTP Protocol (revision 2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1534, Interoperation between DHCP and BOOTP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1542, Clarifications and Extensions for the Bootstrap Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2131, Dynamic Host Configuration Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2132, DHCP Options and BOOTP Vendor Extensions is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2347, TFTP Option Extension is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2348, TFTP Blocksize Option is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2349, TFTP Timeout Interval and Transfer Size Options is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2428, FTP Extensions for IPv6 and NATs is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2865, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2866, RADIUS Accounting is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3046, DHCP Relay Agent Information Option (Option 82) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3596, DNS Extensions to Support IP version 6 is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3768, Virtual Router Redundancy Protocol (VRRP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4251, The Secure Shell (SSH) Protocol Architecture is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4252, The Secure Shell (SSH) Authentication Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

All 7210 platforms support password and publickey based user authentication. 7210 SAS-D support only password based authentication.

RFC 4253, The Secure Shell (SSH) Transport Layer Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4254, The Secure Shell (SSH) Connection Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4632, Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5880, Bidirectional Forwarding Detection (BFD) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5881, Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 6528, Defending against Sequence Number Attacks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7130, Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

13.6 IP — Multicast

RFC 1112, Host Extensions for IP Multicasting is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IGMP v1, v2, v3 is supported.

RFC 2236, Internet Group Management Protocol, Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3306, Unicast-Prefix-based IPv6 Multicast Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3376, Internet Group Management Protocol, Version 3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3446, Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

MSDP supported only on Sx/S-1/10GE standalone.

RFC 3618, Multicast Source Discovery Protocol (MSDP) is supported on Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4604, Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

MLD not supported.

RFC 4607, Source-Specific Multicast for IP is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4608, Source-Specific Protocol Independent Multicast in 232/8 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5059, Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5384, The Protocol Independent Multicast (PIM) Join Attribute Format is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6513, Multicast in MPLS/BGP IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6514, BGP Encodings and Procedures for Multicast in MPLS/IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6515, IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6625, Wildcards in Multicast VPN Auto-Discover Routes is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6826, Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7246, Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7385, IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

13.7 IP — Version 4

RFC 791, Internet Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 792, Internet Control Message Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 826, An Ethernet Address Resolution Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1519, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1812, Requirements for IPv4 Routers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1981, Path MTU Discovery for IP version 6 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2401, Security Architecture for Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Supported only for OSPFv3 authentication. Not supported for services.

RFC 2460, Internet Protocol, Version 6 (IPv6) Specification is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

13.8 IP — Version 6

RFC 2464, Transmission of IPv6 Packets over Ethernet Networks is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 3021, Using 31-Bit Prefixes on IPv4 Point-to-Point Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3122, Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification is supported on K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3587, IPv6 Global Unicast Address Format is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4007, IPv6 Scoped Address Architecture is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4193, Unique Local IPv6 Unicast Addresses is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4291, Internet Protocol Version 6 (IPv6) Addressing Architecture is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4861, Neighbor Discovery for IP version 6 (IPv6) is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5095, Deprecation of Type 0 Routing Headers in IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5952, A Recommendation for IPv6 Address Text Representation is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 6106, IPv6 Router Advertisement Options for DNS Configuration is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6164, Using 127-Bit IPv6 Prefixes on Inter-Router Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

13.9 IPsec

RFC 2401, Security Architecture for the Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only for use with OSPFv3 authentication. Not supported for services.

RFC 2406, IP Encapsulating Security Payload (ESP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only for use with OSPFv3 authentication. Not supported for services.

13.10 IS-IS

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-kaplan-isis-ext-eth-02, Extended Ethernet Frame Size Support is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ISO/IEC 10589:2002, Second Edition, Nov. 2002, Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3359, Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3719, Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3787, Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4971, Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5120, M-ISIS: Multi Topology (MT) Routing in IS-IS is supported on K12, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5130, A Policy Control Mechanism in IS-IS Using Administrative Tags is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5301, Dynamic Hostname Exchange Mechanism for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5302, Domain-wide Prefix Distribution with Two-Level IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5304, IS-IS Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5305, IS-IS Extensions for Traffic Engineering TE is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5306, Restart Signaling for IS-IS (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5308, Routing IPv6 with IS-IS is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5310, IS-IS Generic Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6232, Purge Originator Identification TLV for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6233, IS-IS Registry Extension for Purges is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

K12, K30 support only a single instance and can operate in multi-instance deployment as it supports the processing of TLVs for multi-instance support.

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

13.11 Management

draft-ietf-snmpv3-update-mib-05, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-idr-bgp4-mib-05, Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-wg-mib-06, Management Information Base for Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mboned-msdp-mib-01, Multicast Source Discovery protocol MIB is supported on Sx/S-1/10GE



Note:

Only in standalone mode.

draft-ietf-mpls-ldp-mib-07, Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mpls-lsr-mib-06, Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mpls-te-mib-04, Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-mib-update-08, OSPF Version 2 Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaaddressfamilynumbers-mib, IANA-ADDRESS-FAMILY-NUMBERS-MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaifttype-mib, IANAifType-MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaiprouteprotocol-mib, IANA-RTPROTO-MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8021-CFM-MIB, IEEE P802.1ag(TM) CFM MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8021-PAE-MIB, IEEE 802.1X MIB is supported on D, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8023-LAG-MIB, IEEE 802.3ad MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

LLDP-MIB, IEEE P802.1AB(TM) LLDP MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1157, A Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1215, A Convention for Defining Traps for use with the SNMP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1724, RIP Version 2 MIB Extension is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2021, Remote Network Monitoring Management Information Base Version 2 using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2138, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2206, RSVP Management Information Base using SMIv2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2213, Integrated Services Management Information Base using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2571, An Architecture for Describing SNMP Management Frameworks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2573, SNMP Applications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2575, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2578, Structure of Management Information Version 2 (SMIv2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2579, Textual Conventions for SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2787, Definitions of Managed Objects for the Virtual Router Redundancy Protocol is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2819, Remote Network Monitoring Management Information Base is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2856, Textual Conventions for Additional High Capacity Data Types is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2863, The Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2864, The Inverted Stack Table Extension to the Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2933, Internet Group Management Protocol MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3014, Notification Log MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3164, The BSD syslog Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3165, Definitions of Managed Objects for the Delegation of Management Scripts is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3231, Definitions of Managed Objects for Scheduling Management Operations is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3273, Remote Network Monitoring Management Information Base for High Capacity Networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3414, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3416, Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3417, Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3419, Textual Conventions for Transport Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3584, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3593, Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3635, Definitions of Managed Objects for the Ethernet-like Interface Types is supported on D, E, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3826, The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3877, Alarm Management Information Base (MIB) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4001, Textual Conventions for Internet Network Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4022, Management Information Base for the Transmission Control Protocol (TCP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4113, Management Information Base for the User Datagram Protocol (UDP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4220, Traffic Engineering Link Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4292, IP Forwarding Table MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4293, Management Information Base for the Internet Protocol (IP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5101, Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information is supported on Mxp, Sx/S-1/10GE, and R6

RFC 6241, Network Configuration Protocol (NETCONF) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

RFC 6242, Using the NETCONF Protocol over Secure Shell (SSH) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

13.12 MPLS — General

RFC 3031, Multiprotocol Label Switching Architecture is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3032, MPLS Label Stack Encoding is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3443, Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4182, Removing a Restriction on the use of MPLS Explicit NULL is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5332, MPLS Multicast Encapsulations is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

13.13 MPLS — GMPLS

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

13.14 MPLS — LDP

draft-ietf-mpls-ldp-ipv6-15, Updates to LDP for IPv6 is supported on Mxp

draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-v2-00, LDP Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-hello-reduce-04, Targeted LDP Hello Reduction is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3037, LDP Applicability is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5036, LDP Specification is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5443, LDP IGP Synchronization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5561, LDP Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:

P2MP LSPs only.

13.15 MPLS — MPLS-TP

RFC 5586, MPLS Generic Associated Channel is supported on T(N), R6, and R12

RFC 5921, A Framework for MPLS in Transport Networks is supported on T(N), R6, and R12

RFC 5960, MPLS Transport Profile Data Plane Architecture is supported on T(N), R6, and R12

RFC 6370, MPLS Transport Profile (MPLS-TP) Identifiers is supported on T(N), R6, and R12

RFC 6378, MPLS Transport Profile (MPLS-TP) Linear Protection is supported on T(N), R6, and R12

RFC 6426, MPLS On-Demand Connectivity and Route Tracing is supported on T(N), R6, and R12

RFC 6428, Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile is supported on T(N), R6, and R12

RFC 6478, Pseudowire Status for Static Pseudowires is supported on T(N), R6, and R12

RFC 7213, MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing is supported on T(N), R6, and R12

13.16 MPLS — OAM

RFC 6424, Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6425, Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping is supported on T(N), Mxp, R6, and R12

13.17 MPLS — RSVP-TE

RFC 2702, Requirements for Traffic Engineering over MPLS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2747, RSVP Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2961, RSVP Refresh Overhead Reduction Extensions is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3097, RSVP Cryptographic Authentication -- Updated Message Type Value is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3477, Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4561, Definition of a Record Route Object (RRO) Node-Id Sub-Object is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4875, Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs) is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4950, ICMP Extensions for Multiprotocol Label Switching is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5817, Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

13.18 OSPF

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1765, OSPF Database Overflow is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2328, OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3101, The OSPF Not-So-Stubby Area (NSSA) Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3509, Alternative Implementations of OSPF Area Border Routers is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3623, Graceful OSPF Restart Graceful OSPF Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3630, Traffic Engineering (TE) Extensions to OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4222, Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4552, Authentication/Confidentiality for OSPFv3 is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4576, Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4970, Extensions to OSPF for Advertising Optional Router Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5185, OSPF Multi-Area Adjacency is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5187, OSPFv3 Graceful Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5243, OSPF Database Exchange Summary List Optimization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5250, The OSPF Opaque LSA Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5340, OSPF for IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5838, Support of Address Families in OSPFv3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6987, OSPF Stub Router Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

13.19 Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, OAM Procedures for VPWS Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3916, Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3985, Pseudo Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4385, Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4446, IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4447, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5659, An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6073, Segmented Pseudowire is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6310, Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6391, Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network is supported on K12, K30, Mxp, R6, and R12

RFC 6718, Pseudowire Redundancy is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6870, Pseudowire Preferential Forwarding Status bit is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7023, MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7267, Dynamic Placement of Multi-Segment Pseudowires is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

13.20 Quality of Service

RFC 2430, A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2598, An Expedited Forwarding PHB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3140, Per Hop Behavior Identification Codes is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3260, New Terminology and Clarifications for Diffserv is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

13.21 RIP

RFC 1058, Routing Information Protocol is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2082, RIP-2 MD5 Authentication is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2453, RIP Version 2 is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

13.22 Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, IEEE default profile is supported only includes the Dxp-12p ETR, Dxp-16p, Dxp-24p. Sx-10/100GE does not support IEEE default profile.

ITU-T G.781, Synchronization layer functions, issued 09/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8261, Timing and synchronization aspects in packet networks, issued 04/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8262, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8264, Distribution of timing information through packet networks, issued 10/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8265.1, Precision time protocol telecom profile for frequency synchronization, issued 10/2010 is supported on D-ETR, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, R6, and R12

ITU-T G.8275.1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014 is supported on Dxp, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12



Note:

For 7210 SAS-Sx 10/100GE, the support only includes the Sx 10/100GE QSFP28 variant. For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

IEC/IEEE 61850-9-3-2016, Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation is supported on Dxp-16p and Dxp-24p

IEEE C37.238-2017 - IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications is supported on Dxp-16p and Dxp-24p

RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

13.23 VPLS

RFC 4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:


On 7210 platforms, only BGP-AD is supported with TLDP signalling for PW. No BGP signalling is supported for PW establishment.

RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5501, Requirements for Multicast Support in Virtual Private LAN Services is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)