



7705 Service Aggregation Router

Release 26.4.R1

Quality of Service Guide

3HE 24958 AAAA TQZZA

Edition: 01

April 2026

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

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

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

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

No part of this document may be modified.

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

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

© 2026 Nokia.

Table of contents

List of tables.....	14
List of figures.....	18
1 Preface.....	20
1.1 Audience.....	20
1.2 Technical support.....	20
2 7705 SAR QoS configuration process.....	21
3 QoS and QoS policies.....	22
3.1 QoS overview.....	22
3.1.1 Overview.....	23
3.1.2 Egress and ingress traffic direction.....	24
3.1.2.1 Ring traffic.....	25
3.1.3 Forwarding classes.....	26
3.1.4 Scheduling modes.....	27
3.1.5 Intelligent discards.....	27
3.1.6 Buffering.....	29
3.1.6.1 Buffer pools.....	29
3.1.6.2 CBS and MBS configuration.....	29
3.1.6.3 Buffer unit allocation and buffer chaining.....	30
3.1.7 Per-SAP aggregate shapers (H-QoS) on Gen-2 hardware.....	32
3.1.7.1 Shaped and unshaped SAPs.....	32
3.1.7.2 H-QoS example.....	35
3.1.8 Per-VLAN network egress shapers.....	35
3.1.8.1 Shaped and unshaped VLANs.....	36
3.1.8.2 Per-VLAN shaper support.....	37
3.1.8.3 VLAN shaper applications.....	37
3.1.9 Per-customer aggregate shapers (multiservice site) on Gen-2 hardware.....	39
3.1.9.1 MSS support.....	40
3.1.9.2 MSS and LAG interaction on the 7705 SAR-8 Shelf V2 and 7705 SAR-18.....	41
3.1.10 QoS for hybrid ports on Gen-2 hardware.....	42
3.1.11 QoS for Gen-3 adapter cards and platforms.....	45

3.1.11.1	6-port SAR-M Ethernet module.....	46
3.1.11.2	4-priority scheduling behavior on Gen-3 hardware.....	46
3.1.11.3	Gen-3 hardware and LAG.....	54
3.1.12	QoS on a ring adapter card or module.....	54
3.1.12.1	Network and network queue QoS policy types.....	55
3.1.12.2	Network QoS and network queue policies on a ring adapter card or module....	55
3.1.12.3	Considerations for using ring adapter card or module QoS policies.....	56
3.1.13	QoS for IPSec traffic.....	56
3.1.14	QoS for network group encryption traffic.....	56
3.2	Access ingress.....	56
3.2.1	Access ingress traffic classification.....	57
3.2.1.1	Traffic classification types.....	57
3.2.2	Access ingress queues.....	60
3.2.3	Access ingress queuing and scheduling.....	60
3.2.3.1	Profiled (rate-based) scheduling.....	62
3.2.3.2	Queue-type scheduling.....	62
3.2.3.3	4-priority scheduling.....	63
3.2.3.4	4-priority (Gen-3) scheduling.....	65
3.2.3.5	16-priority scheduling.....	65
3.2.3.6	Ingress queuing and scheduling for BMU traffic.....	66
3.2.4	Access ingress per-SAP aggregate shapers (access ingress H-QoS).....	66
3.2.4.1	Access ingress per-SAP shapers arbitration.....	68
3.2.5	Ingress shaping to fabric (access and network).....	70
3.2.5.1	BMU support.....	70
3.2.5.2	LAG SAP support (access only).....	71
3.2.6	Configurable ingress shaping to fabric (access and network).....	71
3.2.7	Fabric shaping on the fixed platforms (access and network).....	75
3.3	Traffic flow across the fabric.....	75
3.4	Network egress.....	75
3.4.1	BMU traffic at network egress.....	76
3.4.2	Network egress queuing aggregation.....	76
3.4.2.1	Network egress per-VLAN queuing.....	76
3.4.3	Network egress scheduling.....	77
3.4.3.1	Network egress 4-priority scheduling.....	78
3.4.3.2	Network egress 4-priority (Gen-3) scheduling.....	79
3.4.3.3	Network egress 16-priority scheduling.....	79

3.4.4	Network egress shaping.....	80
3.4.5	Network egress shaping for hybrid ports.....	80
3.4.6	Network egress per-VLAN shapers.....	80
3.4.6.1	Network egress per-VLAN shapers arbitration.....	82
3.4.7	Network egress marking and re-marking.....	82
3.4.7.1	Network egress marking and re-marking on Ethernet ports.....	82
3.5	Network ingress.....	83
3.5.1	Network ingress classification.....	83
3.5.1.1	Network ingress tunnel QoS override.....	83
3.5.2	Network ingress queuing.....	84
3.5.2.1	Network ingress queuing for BMU traffic.....	85
3.5.3	Network ingress scheduling.....	85
3.5.3.1	Network ingress 4-priority scheduling.....	86
3.5.3.2	Network ingress 4-priority (Gen-3) scheduling.....	87
3.5.3.3	Network ingress 16-priority scheduling.....	88
3.5.4	Network ingress shaping to fabric.....	88
3.5.5	Configurable network ingress shaping to fabric.....	88
3.5.6	Network fabric shaping on the fixed platforms.....	89
3.6	Access egress.....	89
3.6.1	Access egress queuing and scheduling.....	89
3.6.1.1	BMU traffic access egress queuing and scheduling.....	91
3.6.1.2	ATM access egress queuing and scheduling.....	91
3.6.1.3	Ethernet access egress queuing and scheduling.....	93
3.6.2	Access egress per-SAP aggregate shapers (access egress H-QoS).....	94
3.6.2.1	Access egress per-SAP shapers arbitration.....	94
3.6.3	Access egress shaping for hybrid ports.....	95
3.6.4	Access egress for 4-priority (Gen-3) scheduling.....	95
3.6.5	Access egress marking and re-marking.....	95
3.6.6	Packet byte offset.....	96
3.7	QoS policies overview.....	98
3.7.1	Overview.....	99
3.7.2	Service ingress QoS policies.....	101
3.7.3	Service egress QoS policies.....	103
3.7.4	MC-MLPPP SAP egress QoS policies.....	105
3.7.5	Network and network queue QoS policies.....	106
3.7.5.1	Network QoS policies.....	106

3.7.5.2	Network queue QoS policies.....	113
3.7.6	Network and service QoS queue parameters.....	115
3.7.6.1	Queue ID.....	115
3.7.6.2	Committed information rate.....	115
3.7.6.3	Peak information rate.....	116
3.7.6.4	Adaptation rule.....	116
3.7.6.5	Committed burst size.....	117
3.7.6.6	Maximum burst size.....	117
3.7.6.7	High-priority-only buffers.....	117
3.7.6.8	High and low enqueueing thresholds.....	118
3.7.6.9	Queue counters.....	118
3.7.6.10	Queue type.....	119
3.7.6.11	Queue mode.....	119
3.7.6.12	Rate limiting.....	120
3.7.7	Slope policies (WRED and RED).....	121
3.7.7.1	WRED MinThreshold and MaxThreshold computation.....	122
3.7.7.2	WRED on bridging domain (ring) queues.....	122
3.7.7.3	Payload-based WRED.....	123
3.7.8	ATM traffic descriptor profiles.....	124
3.7.9	Fabric profiles.....	124
3.7.10	Shaper policies.....	124
3.7.11	QoS policy entities.....	124
3.8	Configuration notes.....	125
4	Network QoS policies.....	126
4.1	Overview.....	126
4.2	Basic configuration.....	126
4.2.1	Configuring a network QoS policy.....	127
4.2.2	Creating a network QoS policy.....	128
4.2.3	Applying network QoS policies.....	129
4.2.4	Default network QoS policy values.....	130
4.3	Service management tasks.....	132
4.3.1	Deleting QoS policies.....	132
4.3.2	Copying and overwriting network policies.....	133
4.3.3	Editing QoS policies.....	133
4.4	Network QoS policy command reference.....	134

4.4.1	Command hierarchies.....	134
4.4.1.1	Configuration commands.....	134
4.4.1.2	Operational commands.....	135
4.4.1.3	Show commands.....	135
4.4.2	Command descriptions.....	136
4.4.2.1	Configuration commands.....	136
4.4.2.2	Operational commands.....	157
4.4.2.3	Show commands.....	158
5	Network queue QoS policies.....	171
5.1	Overview.....	171
5.2	Basic configuration.....	171
5.2.1	Configuring a network queue QoS policy.....	172
5.2.2	Creating a network queue QoS policy.....	172
5.2.3	Applying network queue QoS policies.....	173
5.2.3.1	Adapter cards.....	173
5.2.3.2	Network ports.....	175
5.2.4	Configuring per-VLAN network egress shapers.....	176
5.2.5	Configuring a CIR for network egress unshaped VLANs.....	177
5.2.6	Default network queue QoS policy values.....	177
5.3	Service management tasks.....	181
5.3.1	Deleting QoS policies.....	181
5.3.2	Copying and overwriting QoS policies.....	182
5.3.3	Editing QoS policies.....	183
5.4	Network queue QoS policy command reference.....	184
5.4.1	Command hierarchies.....	184
5.4.1.1	Configuration commands.....	184
5.4.1.2	Operational commands.....	185
5.4.1.3	Show commands.....	185
5.4.2	Command descriptions.....	186
5.4.2.1	Configuration commands.....	186
5.4.2.2	Operational commands.....	202
5.4.2.3	Show commands.....	202
6	Service egress and ingress QoS policies.....	206
6.1	Overview.....	206

6.2	Basic configuration.....	206
6.2.1	Creating service egress and ingress QoS policies.....	207
6.2.1.1	Creating a service egress QoS policy.....	207
6.2.1.2	Creating a service ingress QoS policy.....	210
6.2.1.3	Creating an MC-MLPPP SAP egress QoS policy.....	213
6.2.2	Applying service egress and ingress policies.....	216
6.2.2.1	VLL and VPLS services.....	216
6.2.2.2	IES and VPRN services.....	217
6.2.3	Default service egress and ingress policy values.....	217
6.2.3.1	Service egress policy defaults.....	218
6.2.3.2	Service ingress policy defaults.....	219
6.3	Service management tasks.....	220
6.3.1	Deleting QoS policies.....	220
6.3.1.1	Removing a QoS policy from a service SAP.....	220
6.3.1.2	Removing a policy from the QoS configuration.....	221
6.3.2	Copying and overwriting QoS policies.....	221
6.3.3	Editing QoS policies.....	222
6.4	Service egress and ingress QoS policy command reference.....	223
6.4.1	Command hierarchies.....	223
6.4.1.1	Service egress QoS policy configuration commands.....	223
6.4.1.2	Service ingress QoS policy configuration commands.....	223
6.4.1.3	MC-MLPPP SAP egress QoS policies.....	224
6.4.1.4	Operational commands.....	225
6.4.1.5	Show commands.....	225
6.4.2	Command descriptions.....	226
6.4.2.1	Configuration commands.....	226
6.4.2.2	Operational commands.....	265
6.4.2.3	Show commands.....	266
7	Slope QoS policies.....	297
7.1	Overview.....	297
7.2	Basic configuration.....	297
7.2.1	Creating a slope QoS policy.....	298
7.2.2	Applying slope policies.....	299
7.2.3	Default slope policy values.....	299
7.3	Service management tasks.....	300

7.3.1	Deleting QoS policies.....	300
7.3.1.1	Removing a policy from the QoS configuration.....	300
7.3.2	Copying and overwriting QoS policies.....	301
7.3.3	Editing QoS policies.....	302
7.4	Slope QoS policy command reference.....	303
7.4.1	Command hierarchies.....	303
7.4.1.1	Configuration commands.....	303
7.4.1.2	Operational commands.....	303
7.4.1.3	Show commands.....	303
7.4.2	Command descriptions.....	304
7.4.2.1	Configuration commands.....	304
7.4.2.2	Operational commands.....	309
7.4.2.3	Show commands.....	309
8	ATM QoS traffic descriptor profiles.....	313
8.1	ATM traffic descriptor profiles.....	313
8.1.1	ATM traffic management.....	313
8.1.1.1	ATM service categories.....	313
8.1.1.2	ATM traffic descriptors and QoS parameters.....	314
8.1.1.3	ATM policing.....	314
8.1.1.4	Shaping.....	315
8.1.1.5	ATM queuing and scheduling.....	315
8.1.1.6	Congestion avoidance.....	315
8.2	Basic configuration.....	316
8.2.1	Creating an ATM traffic descriptor profile QoS policy.....	316
8.2.2	Applying ATM traffic descriptor profile policies.....	317
8.2.2.1	ATM VLL (Apipe) SAPs.....	317
8.2.3	Default ATM traffic descriptor profile policy values.....	317
8.3	Service management tasks.....	318
8.3.1	Removing an ATM traffic descriptor profile from the QoS configuration.....	318
8.3.2	Copying and overwriting an ATM traffic descriptor profile.....	319
8.3.3	Editing QoS policies.....	319
8.4	ATM QoS policy command reference.....	320
8.4.1	Command hierarchies.....	320
8.4.1.1	Configuration commands.....	320
8.4.1.2	Operational commands.....	320

8.4.1.3	Show commands.....	320
8.4.2	Command descriptions.....	321
8.4.2.1	Configuration commands.....	321
8.4.2.2	Operational commands.....	329
8.4.2.3	Show commands.....	330
9	QoS fabric profiles.....	335
9.1	Basic configuration.....	335
9.1.1	Creating a QoS fabric profile.....	335
9.1.2	Applying a QoS fabric profile.....	337
9.1.3	Default fabric profile values.....	337
9.2	Service management tasks.....	338
9.2.1	Removing a fabric profile from the QoS configuration.....	338
9.2.2	Copying and overwriting a fabric profile.....	338
9.2.3	Editing QoS policies.....	338
9.3	QoS fabric profile command reference.....	339
9.3.1	Command hierarchies.....	339
9.3.1.1	Configuration commands.....	339
9.3.1.2	Operational commands.....	339
9.3.1.3	Show commands.....	339
9.3.2	Command descriptions.....	341
9.3.2.1	Configuration commands.....	341
9.3.2.2	Operational commands.....	348
9.3.2.3	Show commands.....	348
10	QoS shapers and shaper QoS policies.....	353
10.1	Overview.....	353
10.2	Basic configuration.....	353
10.2.1	Creating a shaper QoS policy and shaper groups.....	354
10.2.2	Applying a shaper QoS policy and shaper groups.....	355
10.2.2.1	Applying a shaper policy.....	356
10.2.2.2	Applying a shaper group.....	357
10.2.2.3	Viewing shaper policy information.....	357
10.2.3	Default shaper QoS policy values.....	358
10.2.4	Configuring per-SAP aggregate shapers and an unshaped SAP aggregate shaper (H-QoS).....	358

10.2.4.1	Creating 16-priority shaped SAPs and configuring per-SAP aggregate shapers.....	359
10.2.4.2	Configuring an unshaped aggregate CIR for all 4-priority unshaped SAPs (access ingress).....	361
10.2.4.3	Configuring an unshaped aggregate CIR for all 4-priority unshaped SAPs (access egress).....	362
10.2.5	Configuring per-VLAN shapers and an unshaped VLAN shaper.....	362
10.2.5.1	Configuring per-VLAN network egress shapers.....	363
10.2.5.2	Configuring a CIR for network egress unshaped VLANs.....	363
10.3	Service management tasks.....	363
10.3.1	Removing and deleting QoS policies.....	363
10.3.2	Copying and overwriting QoS policies.....	364
10.3.3	Editing QoS policies.....	365
10.4	Shaper QoS policy command reference.....	366
10.4.1	Command hierarchies.....	366
10.4.1.1	Configuration commands.....	366
10.4.1.2	Operational commands.....	366
10.4.1.3	Show commands.....	366
10.4.2	Command descriptions.....	367
10.4.2.1	Configuration commands.....	367
10.4.2.2	Operational commands.....	371
10.4.2.3	Show commands.....	371
11	Security QoS and security QoS policies.....	374
11.1	Overview.....	374
11.2	QoS for firewall-extracted packets to the CSM.....	374
11.3	Multi-chassis firewall QoS.....	375
11.4	Security queue QoS policies.....	375
11.4.1	Packet queuing with DSCP.....	375
11.5	Basic configuration.....	376
11.5.1	Creating a security data queue QoS policy.....	376
11.5.2	Default security queue policy parameter values.....	377
11.6	Service management tasks.....	377
11.6.1	Deleting QoS policies.....	378
11.6.2	Copying and overwriting QoS policies.....	378
11.6.3	Editing QoS policies.....	378
11.7	Security queue QoS policy command reference.....	379

11.7.1	Command hierarchies.....	379
11.7.1.1	Configuration commands.....	379
11.7.1.2	Operational commands.....	379
11.7.1.3	Show commands.....	379
11.7.2	Command descriptions.....	380
11.7.2.1	Configuration commands.....	380
11.7.2.2	Operational commands.....	384
11.7.2.3	Show commands.....	385
12	List of acronyms.....	388
13	Supported standards and protocols.....	415
13.1	Telecom standards.....	415
13.2	Protocol support.....	416
13.2.1	ATM.....	416
13.2.2	BFD.....	416
13.2.3	BGP.....	416
13.2.4	DHCP/DHCPv6.....	417
13.2.5	Differentiated services.....	417
13.2.6	Digital data network management.....	418
13.2.7	ECMP.....	418
13.2.8	Ethernet VPN (EVPN).....	418
13.2.9	Frame relay.....	418
13.2.10	GRE.....	418
13.2.11	Internet protocol (IP) – version 4.....	418
13.2.12	Internet protocol (IP) – version 6.....	419
13.2.13	IPSec.....	419
13.2.14	IS-IS.....	420
13.2.15	LDP.....	420
13.2.16	LDP and IP FRR.....	421
13.2.17	MPLS.....	421
13.2.18	MPLS – OAM.....	422
13.2.19	Multicast.....	422
13.2.20	Network management.....	422
13.2.21	OSPF.....	423
13.2.22	OSPFv3.....	424

13.2.23	PPP.....	424
13.2.24	Pseudowires.....	424
13.2.25	RIP.....	425
13.2.26	RADIUS.....	425
13.2.27	RSVP-TE and FRR.....	425
13.2.28	Segment routing (SR).....	426
13.2.29	SONET/SDH.....	426
13.2.30	SSH.....	426
13.2.31	Synchronization.....	426
13.2.32	TACACS+.....	427
13.2.33	TLS.....	427
13.2.34	TWAMP.....	427
13.2.35	VPLS.....	427
13.2.36	VRRP.....	428
13.3	Proprietary MIBs.....	428

List of tables

Table 1: Configuration process.....	21
Table 2: Default forwarding classes.....	26
Table 3: Ingress packet priority during overload.....	28
Table 4: Buffer support on adapter cards and platforms.....	30
Table 5: Scheduling behavior for Gen-2 and Gen-3 hardware.....	45
Table 6: Traffic classification types.....	57
Table 7: Access ingress traffic classification for SAPs per service type	59
Table 8: Scheduling modes supported by adapter cards and ports at access ingress.....	61
Table 9: Fabric profile mode options and capabilities.....	72
Table 10: Scheduling modes supported by adapter cards and ports at network egress.....	77
Table 11: Default network ingress QoS policy.....	84
Table 12: Scheduling modes supported by adapter cards and ports at network ingress.....	85
Table 13: Scheduling modes supported by adapter cards and ports at access egress.....	89
Table 14: Scheduler weight values (WRR) based on MIR for T1/E1 ASAP Adapter cards and 2-port OC3/STM1 Channelized Adapter card.....	92
Table 15: Scheduler weight values (WRR) based on MIR for the 4-port OC3/STM1 Clear Channel Adapter card.....	92
Table 16: ATM scheduling and relative priorities.....	93
Table 17: PBO for SAPs and platforms	96
Table 18: QoS policy types and descriptions.....	100
Table 19: Forwarding class and enqueueing priority classification hierarchy based on rule type.....	102
Table 20: Default service ingress policy ID 1 definition.....	103

Table 21: Default service egress policy ID 1 definition.....	104
Table 22: MC-MLPPP class priorities.....	105
Table 23: Packet forwarding class to MLPPP class mapping.....	105
Table 24: Default network QoS policy egress marking.....	107
Table 25: Default network QoS policy DSCP-to-forwarding class mappings.....	108
Table 26: Default network QoS policy LSP EXP-to-forwarding class mappings.....	109
Table 27: Default network QoS policy dot1p-to-queue class mappings.....	110
Table 28: Applications and support for configurable DSCP or dot1p markings.....	111
Table 29: Default network queue policy definition.....	113
Table 30: Payload-based WRED: discards and tail drop starts.....	123
Table 31: Valid DSCP names.....	142
Table 32: DSCP-to-default FC value mapping	154
Table 33: DSCP name-to-value mapping field descriptions.....	160
Table 34: Network policy field descriptions.....	163
Table 35: Ethernet ring network policy field descriptions.....	165
Table 36: Application QoS field descriptions.....	168
Table 37: DSCP-to-FC mapping field descriptions.....	170
Table 38: Default network queue policy definitions.....	177
Table 39: CBS forwarding class defaults.....	192
Table 40: High-prio-only forwarding class defaults.....	194
Table 41: MBS forwarding class defaults.....	195
Table 42: Network queue policy field descriptions.....	204
Table 43: Service egress policy defaults.....	218

Table 44: Service ingress policy defaults.....	219
Table 45: Valid DSCP names.....	235
Table 46: SAP egress field descriptions.....	269
Table 47: SAP ingress field descriptions.....	275
Table 48: Buffer pool field descriptions.....	294
Table 49: Slope policy defaults.....	299
Table 50: Slope policy field descriptions.....	311
Table 51: ATM traffic descriptors.....	314
Table 52: ATM TD profile defaults.....	317
Table 53: Service category descriptor type default values.....	323
Table 54: Traffic descriptor type parameters.....	323
Table 55: ATM service categories.....	325
Table 56: Default shaping values.....	326
Table 57: Service category traffic descriptor parameters.....	327
Table 58: ATM traffic parameter defaults.....	328
Table 59: ATM traffic descriptor profile field descriptions.....	331
Table 60: SAP field descriptions.....	333
Table 61: Fabric profile defaults.....	337
Table 62: QoS fabric profile field descriptions.....	350
Table 63: System QoS field descriptions.....	351
Table 64: Shaper policy defaults.....	358
Table 65: Shaper policy defaults.....	368
Table 66: Shaper policy field descriptions.....	373

Table 67: Security queue parameter defaults.....	377
Table 68: Security policy field descriptions.....	387
Table 69: Acronyms.....	388

List of figures

Figure 1: Egress and ingress traffic direction.....	25
Figure 2: Ingress and egress traffic on a 2-port 10GigE (Ethernet) Adapter card.....	26
Figure 3: Access ingress scheduling for 4-priority and 16-priority SAPs (with per-SAP aggregate shapers).....	33
Figure 4: Access egress scheduling for 4-priority and 16-priority SAPs (with per-SAP aggregate shapers) (per port).....	34
Figure 5: Network egress shaped and unshaped VLAN queuing and scheduling.....	36
Figure 6: VLAN shapers for dual uplinks.....	38
Figure 7: VLAN shapers in aggregation site scenario.....	39
Figure 8: Hybrid port egress shapers and schedulers on Gen-2 hardware.....	44
Figure 9: 4-priority scheduling at access ingress (Gen-3 hardware).....	48
Figure 10: 4-priority scheduling at access egress (Gen-3 hardware).....	49
Figure 11: 4-priority scheduling at network ingress (Gen-3 hardware): per-destination mode.....	50
Figure 12: 4-priority scheduling at network ingress (Gen-3 hardware): aggregate mode.....	51
Figure 13: 4-priority scheduling at network egress (Gen-3 hardware) on a network port.....	52
Figure 14: 4-priority scheduling for hybrid port egress (Gen-3 hardware).....	53
Figure 15: Ports on a 2-port 10GigE (Ethernet) Adapter card.....	55
Figure 16: 4-priority scheduling.....	65
Figure 17: Access ingress scheduling for 4-priority and 16-priority SAPs (with per-SAP aggregate shapers).....	67
Figure 18: Access ingress per-SAP arbitration to fabric.....	69
Figure 19: Fabric shapers in per-destination mode.....	73
Figure 20: Fabric shapers in aggregate mode.....	74

Figure 21: Network egress shaped and unshaped VLAN queuing and scheduling.....	81
Figure 22: Access egress 16-priority and 4-priority per-SAP arbitration for a single port.....	95
Figure 23: Traffic queuing model for three queues and three classes.....	102
Figure 24: WRED for high-priority and low-priority traffic in the same queue.....	122

1 Preface

This guide describes the Quality of Service (QoS) functionality provided by the 7705 SAR and presents configuration and implementation examples.

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



Note: This manual generically covers Release 26.x content and may contain some content that will be released in later maintenance loads. See the 7705 SAR 26.x.Rx Software Release Notes, part number 3HE24971000xTQZZA, for information about features supported in each load of the Release 26.x software.



Note: As of Release 23.4, software support for the following hardware has been deprecated:

- 8-port Ethernet Adapter card, version 2 (a8-ethv2) (3HE02776)
- 12-port Serial Data Interface card, version 1 (a12-sdi) (3HE03391)
- 7705 SAR-W (3HE07349)

These components are no longer recognized in the release.

If information about any of the above components is required, please see the applicable installation guides in Release 22.10.

1.1 Audience

This guide is intended for network administrators who are responsible for configuring the 7705 SAR routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Concepts described in this guide include the following:

- CLI concepts
- QoS policies and profiles

1.2 Technical support

If you purchased a service agreement for your 7705 SAR router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased a Nokia service agreement, follow this link to contact a Nokia support representative and to access product manuals and documentation updates:

[Product Support Portal](#)

2 7705 SAR QoS configuration process

The following table lists the tasks that are required to configure and apply Quality of Service (QoS) policies.

This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 1: Configuration process

Area	Task/description	Chapter
Overview information	General information about QoS policies	QoS and QoS policies
Service configuration	Configure QoS policies:	
	Network	Network QoS policies
	Network queue	Network queue QoS policies
	Service egress/service ingress	Service egress and ingress QoS policies
	Slope	Slope QoS policies
	ATM traffic descriptor	ATM QoS traffic descriptor profiles
	Fabric profile	QoS fabric profiles
	Shaper	QoS shapers and shaper QoS policies
	Security queue	Security QoS and security QoS policies
Reference	List of security and telecom standards, supported protocols, and proprietary MIBs	Supported standards and protocols

3 QoS and QoS policies

This chapter provides an overview of the 7705 SAR Quality of Service (QoS) and information about QoS policy management.

Topics in this chapter include:

- [QoS overview](#)
- [Access ingress](#)
- [Traffic flow across the fabric](#)
- [Network egress](#)
- [Network ingress](#)
- [Access egress](#)
- [QoS policies overview](#)
- [Configuration notes](#)

3.1 QoS overview

This section contains the following overview topics related to QoS:

- [Overview](#)
- [Egress and ingress traffic direction](#)
- [Forwarding classes](#)
- [Scheduling modes](#)
- [Intelligent discards](#)
- [Buffering](#)
- [Per-SAP aggregate shapers \(H-QoS\) on Gen-2 hardware](#)
- [Per-VLAN network egress shapers](#)
- [Per-customer aggregate shapers \(multiservice site\) on Gen-2 hardware](#)
- [QoS for hybrid ports on Gen-2 hardware](#)
- [QoS for Gen-3 adapter cards and platforms](#)
- [QoS on a ring adapter card or module](#)
- [QoS for IPSec traffic](#)
- [QoS for network group encryption traffic](#)

3.1.1 Overview

To provide what network engineers call Quality of Service (QoS), the flow of data in the form of packets must be predetermined and resources must be somehow assured for that predetermined flow. Simple routing does not provide a predetermined path for the traffic, and priorities that are described by Class of Service (CoS) coding simply increase the odds of successful transit for one packet over another. There is still no guarantee of service quality. The guarantee of service quality is what distinguishes QoS from CoS. CoS is an element of overall QoS.

By using the traffic management features of the 7705 SAR, network engineers can achieve a QoS for their customers. Multiprotocol label switching (MPLS) provides a predetermined path, while policing, shaping, scheduling, and marking features ensure that traffic flows in a predetermined and predictable manner.

There is a need to distinguish between high-priority (that is, mission-critical traffic like signaling) and best-effort traffic priority levels when managing traffic flow. Within these priority levels, it is important to have a second level of prioritization, that is, between a certain volume of traffic that is contracted/needed to be transported, and the amount of traffic that is transported if the system resources allow. Throughout this guide, contracted traffic is referred to as in-profile traffic. Traffic that exceeds the user-configured traffic limits is either serviced using a lower priority or discarded in an appropriate manner to ensure that an overall quality of service is achieved.

The 7705 SAR must be properly configured to provide QoS. To ensure end-to-end QoS, each and every intermediate node together with the egress node must be coherently configured. Proper QoS configuration requires careful end-to-end planning, allocation of appropriate resources and coherent configuration among all the nodes along the path of a given service. Once properly configured, each service provided by the 7705 SAR will be contained within QoS boundaries associated with that service and the general QoS parameters assigned to network links.

The 7705 SAR is designed with QoS mechanisms at both egress and ingress to support different customers and different services per physical interface or card, concurrently and harmoniously (see [Egress and ingress traffic direction](#) for a definition of egress and ingress traffic). The 7705 SAR has extensive and flexible capabilities to classify, police, shape and mark traffic to make this happen.



Note: The characteristics and nature of traffic flows in the ingress and egress directions are usually totally different. As an example, traffic is usually shaped at egress for pacing purposes and jitter tolerance imposed by the network transport rules, whereas at ingress, traffic is usually policed to ensure it fits into the traffic volumes defined in the service-level agreement (SLA). Therefore, segregation between ingress and egress offers not only the seamless flexibility to address different requirements but as well allows fine-tuning of appropriate parameters in each direction.

The 7705 SAR supports multiple forwarding classes (FCs) and associated class-based queuing. Ingress traffic can be classified to multiple FCs, and the FCs can be flexibly associated with queues. This provides the ability to control the priority and drop priority of a packet while allowing the fine-tuning of bandwidth allocation to individual flows.

Each forwarding class is important only in relation to the other forwarding classes. A forwarding class allows network elements to weigh the relative importance of one packet over another. With such flexible queuing, packets belonging to a specific flow within a service can be preferentially forwarded based on the CoS of a queue. The forwarding decision is based on the forwarding class of the packet, as assigned by the ingress QoS policy defined for the service access point (SAP).

7705 SAR routers use QoS policies to control how QoS is handled at distinct points in the service delivery model within the device. QoS policies act like a template. Once a policy is created, it can be applied to

many other similar services and ports. As an example, if there is a group of Node Bs connected to a 7705 SAR node, one QoS policy can be applied to all services of the same type, such as High-Speed Downlink Packet Access (HSDPA) offload services.

There are different types of QoS policies that cater to the different QoS needs at each point in the service delivery model. QoS policies are defined in a global context in the 7705 SAR and only take effect when the policy is applied to a relevant entity.

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

The different QoS policies within the 7705 SAR can be divided into two main types.

- QoS policies are used for classification, queue attributes, and marking.
- Slope policies define default buffer allocations and Random Early Discard (RED) and Weighted Random Early Discard (WRED) slope definitions.

The sections that follow provide an overview of the QoS traffic management performed on the 7705 SAR.

3.1.2 Egress and ingress traffic direction

Throughout this document, the terms ingress and egress, when describing traffic direction, are always defined relative to the fabric. For example:

- ingress direction describes packets moving into the switch fabric away from a port (on an adapter card)
- egress direction describes packets moving from the switch fabric and into a port (on an adapter card)

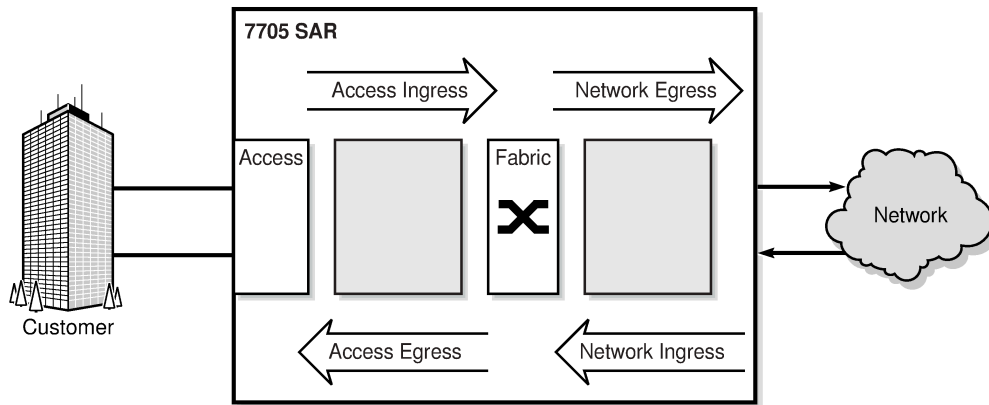
When combined with the terms access and network, which are port and interface modes, the four traffic directions relative to the fabric are (see [Figure 1: Egress and ingress traffic direction](#)):

- access ingress direction describes packets coming in from customer equipment and switched toward the switch fabric
- network egress direction describes packets switched from the switch fabric into the network
- network ingress direction describes packets switched in from the network and moving toward the switch fabric
- access egress direction describes packets switched from the switch fabric toward customer equipment



Note: Throughout this guide, the terms access ingress/egress and service ingress/egress are interchangeable. This section ([QoS overview](#)) uses the term access, and the following sections (beginning with [QoS policies overview](#)) use the term service.

Figure 1: Egress and ingress traffic direction



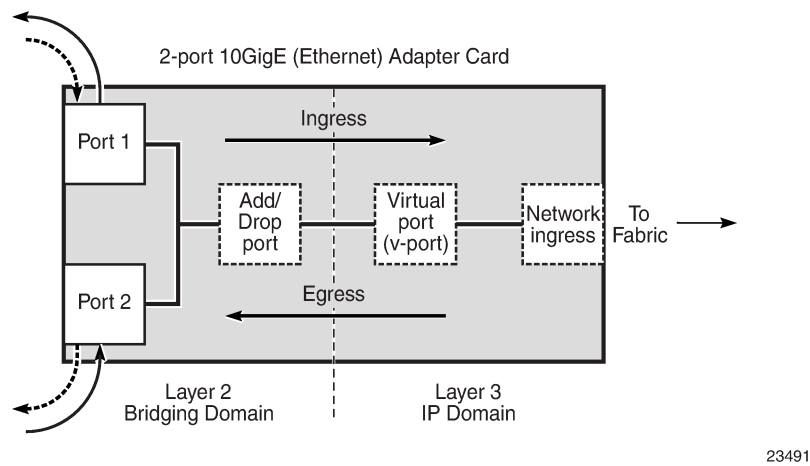
19763

3.1.2.1 Ring traffic

On the 2-port 10GigE (Ethernet) Adapter card and 2-port 10GigE (Ethernet) module, traffic can flow between the Layer 2 bridging domain and the Layer 3 IP domain (see [Figure 2: Ingress and egress traffic on a 2-port 10GigE \(Ethernet\) Adapter card](#)). In the bridging domain, ring traffic flows from one ring port to another, as well as to and from the add/drop port. From the network point of view, traffic from the ring toward the add/drop port and the v-port is considered ingress traffic (drop traffic). Similarly, traffic from the fabric toward the v-port and the add/drop port is considered egress traffic (add traffic).

The 2-port 10GigE (Ethernet) Adapter card or 2-port 10GigE (Ethernet) module functions as an add/drop card to a network side 10 Gb/s optical ring. Conceptually, the card or module should be envisioned as having two domains: a Layer 2 bridging domain where the add/drop function operates and a Layer 3 IP domain where the normal IP processing and IP nodal traffic flows are managed. Ingress and egress traffic flow remains in the context of the nodal fabric. The ring ports are considered to be east-facing and west-facing and are referenced as Port 1 and Port 2. A virtual port (or v-port) provides the interface to the IP domain within the structure of the card or module.

Figure 2: Ingress and egress traffic on a 2-port 10GigE (Ethernet) Adapter card



3.1.3 Forwarding classes

Queues can be created for each forwarding class to determine the manner in which the queue output is scheduled and the type of parameters the queue accepts. The 7705 SAR supports eight forwarding classes per SAP. The following table shows the default mapping of these forwarding classes in order of priority, with Network Control having the highest priority.

Table 2: Default forwarding classes

FC name	FC designation	Queue type	Typical use
Network Control	NC	Expedited	For network control and traffic synchronization
High-1	H1		For delay/jitter sensitive traffic
Expedited	EF		For delay/jitter sensitive traffic
High-2	H2		For delay/jitter sensitive traffic
Low-1	L1	Best Effort	For best-effort traffic
Assured	AF		For best-effort traffic
Low-2	L2		For best-effort traffic
Best Effort	BE		For best-effort traffic

The traffic flows of different forwarding classes are mapped to the queues. This mapping is user-configurable. Each queue has a unique priority. Packets from high-priority queues are scheduled separately, before packets from low-priority queues. More than one forwarding class can be mapped to a single queue. In such a case, the queue type defaults to the priority of the lowest forwarding class (see [Queue type](#) for more information about queue type). By default, the following logical order is followed:

- FC-8 - NC
- FC-7 - H1
- FC-6 - EF
- FC-5 - H2
- FC-4 - L1
- FC-3 - AF
- FC-2 - L2
- FC-1 - BE

At access ingress, traffic can be classified as unicast traffic or one of the multipoint traffic types (broadcast, multicast, or unknown (BMU)). After classification, traffic can be assigned to a queue that is configured to support one of the four traffic types, namely:

- unicast (or implicit)
- broadcast
- multicast
- unknown

3.1.4 Scheduling modes

The scheduler modes available on adapter cards are 4-priority and 16-priority. Which modes are supported on a particular adapter card depends on whether the adapter card is a second-generation or third-generation card.



Note: Throughout the 7705 SAR documentation set, second-generation and third-generation Ethernet adapter cards and Ethernet ports on fixed platforms are also referred to as Gen-2 and Gen-3 hardware.

On Gen-3 hardware, 4-priority scheduling mode is the implicit, default scheduling mode and is not user-configurable. Gen-3 platforms with a TDM block support 4-priority scheduling mode. Gen-2 adapter cards support 16-priority and 4-priority scheduling modes.

For more information about differences between Gen-2 and Gen-3 hardware related to scheduling mode QoS behavior, see [QoS for Gen-3 adapter cards and platforms](#).

For information about scheduling modes as they apply to traffic direction, see the following sections:

- [Access ingress queuing and scheduling](#)
- [Network egress scheduling](#)
- [Network ingress scheduling](#)
- [Access egress queuing and scheduling](#)

3.1.5 Intelligent discards

Most 7705 SAR systems are susceptible to network processor congestion if the packet rate of small packets received on a node or card exceeds the processing capacity. If a node or card receives a high rate of small packet traffic, the node or card enters overload mode. Before the introduction of intelligent

discards, when a node or card entered an overload state, the network processor would randomly drop packets.

The "intelligent discards during overload" feature allows the network processor to discard packets according to a preset priority order. In the egress direction, intelligent discards is applied to traffic entering the card from the fabric.

Traffic is discarded in the following order: low-priority out-of-profile user traffic is discarded first, followed by high-priority out-of-profile user traffic, then low-priority in-profile user traffic, high priority in-profile user traffic, and lastly control plane traffic. In the ingress direction, intelligent discards is applied to traffic entering the card from the physical ports. Traffic is discarded in the following order: low-priority user traffic is always discarded first, followed by high-priority user traffic. This order ensures that low-priority user traffic is always the most susceptible to discards.

In the egress direction, the system differentiates between high-priority and low-priority user traffic based on the internal forwarding class and queue-type fabric header markers. In the ingress direction, the system differentiates between high-priority and low-priority user traffic based on packet header bits. The following table details the classification of user traffic in the ingress direction.

Table 3: Ingress packet priority during overload

Fabric header marker	High-priority values	Low-priority values
MPLS TC	7 to 4	3 to 0
IP DSCP	63 to 32	31 to 0
Eth Dot1p	7 to 4	3 to 0

Intelligent discards during overload ensures priority-based handling of traffic and helps existing traffic management implementations. It does not change how QoS-based classification, buffer management, or scheduling operates on the 7705 SAR. If the node or card is not in overload operation mode, there is no change to the way packets are handled by the network processor.

There are no commands to configure intelligent discards during overload; the feature is automatically enabled on the following cards, modules, and ports:

- 10-port 1GigE/1-port 10GigE X-Adapter card
- 2-port 10GigE (Ethernet) Adapter card (only on the 2.5 Gb/s v-port)
- 2-port 10GigE (Ethernet) module (only on the v-port)
- 8-port Gigabit Ethernet Adapter card
- 6-port Ethernet 10Gbps Adapter card
- Packet Microwave Adapter card
- 4-port SAR-H Fast Ethernet module
- 6-port SAR-M Ethernet module
- 7705 SAR-A Ethernet ports
- 7705 SAR-Ax Ethernet ports
- 7705 SAR-Wx Ethernet ports
- 7705 SAR-M Ethernet ports
- 7705 SAR-H Ethernet ports

- 7705 SAR-Hc Ethernet ports
- 7705 SAR-X Ethernet ports

3.1.6 Buffering

Buffer space is allocated to queues based on the committed buffer space (CBS), the maximum buffer space (MBS) and availability of the resources, and the total amount of buffer space. The CBS and the MBS define the queue depth for a particular queue. The MBS represents the maximum buffer space that is allocated to a particular queue. Whether that much space can actually be allocated depends on buffer usage (that is, the number of other queues and their sizes).

Memory allocation is optimized to guarantee the CBS for each queue. The allocated queue space beyond the CBS is limited by the MBS and depends on the use of buffer space and the guarantees accorded to queues as configured in the CBS.

This section contains information about the following topics:

- [Buffer pools](#)
- [CBS and MBS configuration](#)
- [Buffer unit allocation and buffer chaining](#)

3.1.6.1 Buffer pools

The 7705 SAR supports two types of buffer pools that allocate memory as follows:

- reserved pool – represents the CBS that is guaranteed for all queues. The reserved pool is limited to a maximum of 75% of the total buffer space.
- shared pool – represents the buffer space that remains after the reserved pool has been allocated. The shared pool always has at least 25% of the total buffer space.

Both buffer pools can be displayed in the CLI using the **show pools** command.

3.1.6.2 CBS and MBS configuration

On the access side, CBS is configured in bytes and MBS in bytes or kilobytes using the CLI. See, for example, the **config>qos>sap-ingress/egress>queue>cbs** and **mbs** configuration commands.

On the network side, CBS and MBS values are expressed as a percentage of the total number of available buffers. If the buffer space is further segregated into pools (for example, ingress and egress, access and network, or a combination of these), the CBS and MBS values are expressed as a percentage of the applicable buffer pool. See the **config>qos>network-queue>queue>cbs** and **mbs** configuration commands.

The configured CBS and MBS values are converted to the number of buffers by dividing the CBS or MBS value by a fixed buffer size of 512 bytes or 2304 bytes, depending on the type of adapter card or platform. The number of buffers can be displayed for an adapter card using the **show pools** command.

3.1.6.2.1 Buffer allocation for multicast traffic

When a packet is being multicast to two or more interfaces on the egress adapter card or block of fixed ports, or when a packet at port ingress is mirrored, one extra buffer per packet is used.

In previous releases, this extra buffer was not added to the queue count. When checking CBS during multicast traffic enqueueing, the CBS was divided by two to prevent buffer overconsumption by the extra buffers. As a result, during multicast traffic enqueueing, the CBS buffer limit for the queue was considered reached when half of the available buffers were in use.

As of Release 8.0 of the 7705 SAR, the CBS is no longer divided by two. Instead, the extra buffers are added to the queue count when enqueueing, and are removed from the queue count when the multicast traffic exits the queue. The full CBS value is used, and the extra buffer allocation is visible in buffer allocation displays.

3.1.6.3 Buffer unit allocation and buffer chaining

Packetization buffers and queues are supported in the packet memory of each adapter card or platform. All adapter cards and platforms allocate a fixed space for each buffer. The 7705 SAR supports two buffer sizes: 512 bytes or 2304 bytes, depending on the type of adapter card or platform.

The adapter cards and platforms that support a buffer size of 2304 bytes do not support buffer chaining (see the description below) and only allow a 1-to-1 correspondence of packets to buffers.

The adapter cards and platforms that support a buffer of size of 512 bytes use a method called buffer chaining to process packets that are larger than 512 bytes. To accommodate packets that are larger than 512 bytes, these adapter cards or platforms divide the packet dynamically into a series of concatenated 512-byte buffers. An internal 64-byte header is prepended to the packet, so only 448 bytes of buffer space is available for customer traffic in the first buffer. The remaining customer traffic is split among concatenated 512-byte buffers.

The following table shows the supported buffer sizes on the 7705 SAR adapter cards and platforms. If a version number or variant is not specified, this implies all versions of the adapter card or variants of the platform. Adapter cards and platforms that support buffer chaining have 512 byte buffer size ("Yes"); those that do not support buffer chaining have 2304 byte buffer size ("No").

Table 4: Buffer support on adapter cards and platforms

Adapter card or platform	Buffer space per card/platform (MB)	Buffer chaining support
2-port 10GigE (Ethernet) Adapter card	268 201 (for L2 bridging domain)	Yes Yes (each buffer unit is 768 bytes)
2-port 10GigE (Ethernet) module	201 (for L2 bridging domain)	Yes (each buffer unit is 768 bytes)
2-port OC3/STM1 Channelized Adapter card	310	No
4-port OC3/STM1 / 1-port OC12/STM4 Adapter card	217	Yes

Adapter card or platform	Buffer space per card/platform (MB)	Buffer chaining support
4-port OC3/STM1 Clear Channel Adapter card	352	No
4-port DS3/E3 Adapter card	280	No
6-port E&M Adapter card	38	No
6-port FXS Adapter card	38	No
6-port Ethernet 10Gbps Adapter card	1177	Yes
8-port FXO Adapter card	38	No
8-port Gigabit Ethernet Adapter card	268	Yes
8-port Voice & Teleprotection card	38	No
8-port C37.94 Teleprotection card	38	No
10-port 1GigE/1-port 10GigE X-Adapter card	537	Yes
12-port Serial Data Interface card, version 2 and version 3	268	Yes
16-port T1/E1 ASAP Adapter card	38	No
32-port T1/E1 ASAP Adapter card	57	No
Integrated Services card	268	Yes
Packet Microwave Adapter card	268	Yes
7705 SAR-A	268	Yes
7705 SAR-Ax	268	Yes
7705 SAR-H	268	Yes
7705 SAR-Hc	268	Yes
7705 SAR-M	268	Yes
7705 SAR-Wx	268	Yes
7705 SAR-X (Ethernet ports) ¹	1177	Yes
7705 SAR-X (TDM ports) ¹	46	Yes

Note:

1. The 7705 SAR-X has three buffer pools. Each block of ports (MDA) has its own buffer pool.

3.1.6.3.1 Advantages of buffer chaining

Buffer chaining offers improved efficiency, which is especially evident when smaller packet sizes are transmitted. For example, to queue a 64-byte packet, a card with a fixed buffer of 2304 bytes allocates 2304 bytes, whereas a card with a fixed buffer of 512 bytes allocates only 512 bytes. To queue a 1280-byte packet, a card with a fixed buffer of 2304 bytes allocates 2304 bytes, whereas a card with a fixed buffer of 512 bytes allocates only 1536 bytes (that is, 512 bytes × 3 buffers).

3.1.7 Per-SAP aggregate shapers (H-QoS) on Gen-2 hardware

This section contains overview information as well as information about the following topics:

- [Shaped and unshaped SAPs](#)
- [H-QoS example](#)

This section provides information about per-SAP aggregate shapers for Gen-2 adapter cards and platforms. For information about Gen-3 adapter cards and platforms, see [QoS for Gen-3 adapter cards and platforms](#).

Hierarchical QoS (H-QoS) provides the 7705 SAR with the ability to shape traffic on a per-SAP basis for traffic from up to eight CoS queues associated with that SAP.

On Gen-2 hardware, the per-SAP aggregate shapers apply to access ingress and access egress traffic and operate in addition to the 16-priority scheduler, which must be used for per-SAP aggregate shaping.

The 16-priority scheduler acts as a soft policer, servicing the SAP queues in strict priority order, with conforming traffic (less than CIR) serviced before non-conforming traffic (between CIR and PIR). The 16-priority scheduler on its own cannot enforce a traffic limit on a per-SAP basis; to do this, per-SAP aggregate shapers are required (see [H-QoS example](#)).

The per-SAP shapers are considered aggregate shapers because they shape traffic from the aggregate of one or more CoS queues assigned to the SAP.

[Figure 3: Access ingress scheduling for 4-priority and 16-priority SAPs \(with per-SAP aggregate shapers\)](#) and [Figure 4: Access egress scheduling for 4-priority and 16-priority SAPs \(with per-SAP aggregate shapers\) \(per port\)](#) illustrate per-SAP aggregate shapers for access ingress and access egress, respectively. They indicate how shaped and unshaped SAPs are treated.

H-QoS is not supported on the 4-port SAR-H Fast Ethernet module.

3.1.7.1 Shaped and unshaped SAPs

Shaped SAPs have user-configured rate limits (PIR and CIR)—called the aggregate rate limit—and must use 16-priority scheduling mode. Unshaped SAPs use default rate limits (PIR is maximum and CIR is 0 kb/s) and can use 4-priority or 16-priority scheduling mode.

Shaped 16-priority SAPs are configured with a PIR and a CIR using the **agg-rate-limit** command in the **config>service>service-type service-id>sap** context, where **service-type** is *epipe*, *ipipe*, *ies*, *vprn*, or *vpls* (including routed VPLS). The PIR is set using the *agg-rate* variable and the CIR is set using the *cir-rate* variable.

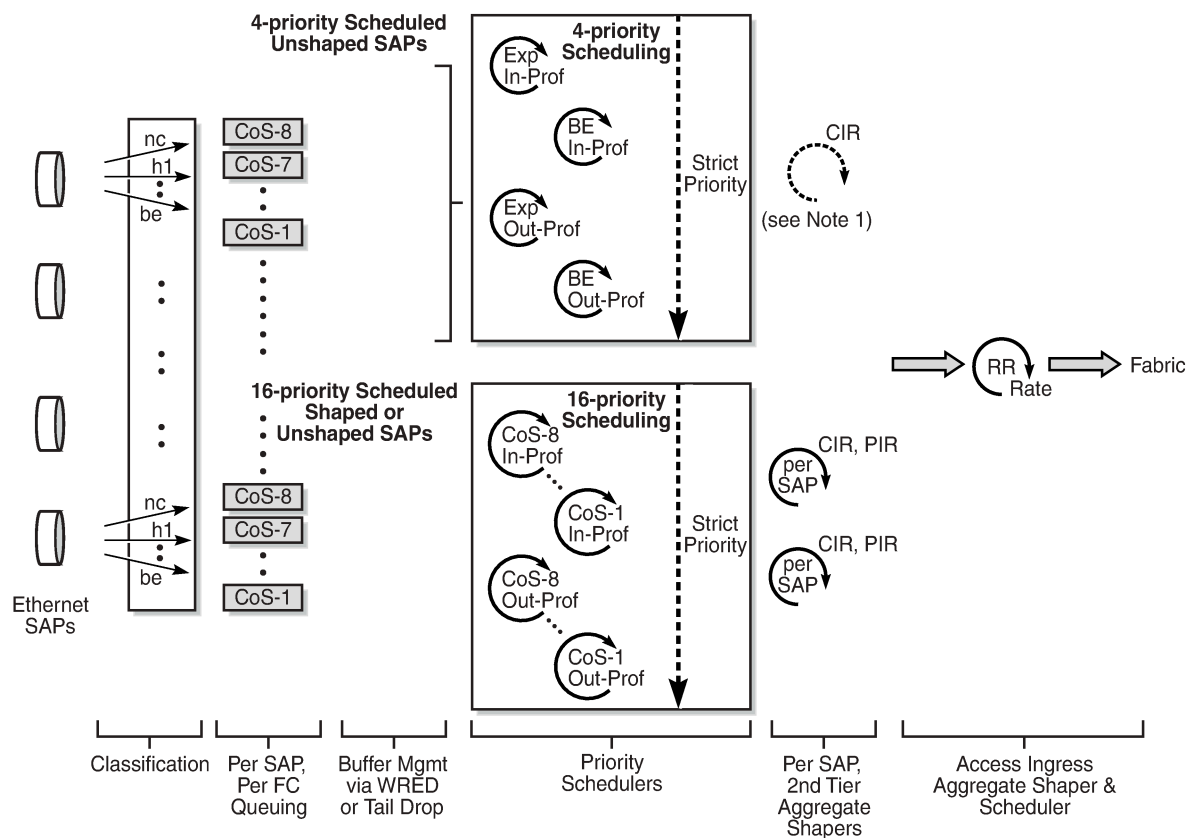
Unshaped 4-priority SAPs are considered unshaped by definition of the default PIR and CIR values (PIR is maximum and CIR is 0 kb/s). Therefore, they do not require any configuration other than to be set to 4-priority scheduling mode.

Unshaped 16-priority SAPs are created when 16-priority scheduling mode is selected, when the default PIR is maximum and the default CIR is 0 kb/s, which are same default settings of a 4-priority SAP. The main reason for preferring unshaped SAPs using 16-priority scheduling over unshaped SAPs using 4-priority scheduling is to have a coherent scheduler behavior (one scheduling model) across all SAPs.

In order for unshaped 4-priority SAPs to compete fairly for bandwidth with 16-priority shaped and unshaped SAPs, a single, aggregate CIR for all the 4-priority SAPs can be configured. This aggregate CIR is applied to all the 4-priority SAPs as a group, not to individual SAPs. In addition, the aggregate CIR is configured differently for access ingress and access egress traffic. On the 7705 SAR-8 Shelf V2 and 7705 SAR-18, access ingress is configured in the **config>qos>fabric-profile** context. On the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-Wx, access ingress is configured in the **config>system>qos>access-ingress-aggregate-rate** context. For all platforms, access egress configuration uses the **config>port>ethernet** context.

For more information about access ingress scheduling and traffic arbitration from the 16-priority and 4-priority schedulers toward the fabric, see [Access ingress per-SAP aggregate shapers \(access ingress H-QoS\)](#).


Figure 3: Access ingress scheduling for 4-priority and 16-priority SAPs (with per-SAP aggregate shapers)



Note 1: Aggregate shaper (CIR) for all the 4-priority unshaped SAPs.

23369

Figure 4: Access egress scheduling for 4-priority and 16-priority SAPs (with per-SAP aggregate shapers) (per port)



Note 1: Aggregate shaper (CIR) for all the 4-priority unshaped SAPs.

23376

3.1.7.1.1 Per-SAP aggregate shaper support

The per-SAP aggregate shapers are supported in both access ingress and access egress directions and can be enabled on the following Ethernet access ports:

- 6-port Ethernet 10Gbps Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card (10-port GigE mode)
- Packet Microwave Adapter card
- 7705 SAR-A
- 7705 SAR-Ax
- 7705 SAR-M
- 7705 SAR-H(all Ethernet access ports except those on the 4-port SAR-H Fast Ethernet module)
- 7705 SAR-Hc
- 7705 SAR-Wx

- 7705 SAR-X

3.1.7.2 H-QoS example

A typical example in which H-QoS is used is where a transport provider uses a 7705 SAR as a PE device and sells 100 Mb/s of fixed bandwidth for point-to-point Internet access, and offers premium treatment to 10% of the traffic. A customer can mark up to 10% of their critical traffic such that it is classified into high-priority queues and serviced before low-priority traffic.

Without H-QoS, there is no way to enforce a limit to ensure that the customer does not exceed the leased 100 Mb/s bandwidth, as illustrated in the following two scenarios:

- If a queue hosting high-priority traffic is serviced at 10 Mb/s and the low-priority queue is serviced at 90 Mb/s, then at a moment when the customer transmits less than 10 Mb/s of high-priority traffic, the customer bandwidth requirement is not met (the transport provider transports less traffic than the contracted rate).
- If the scheduling rate for the high-priority queue is set to 10 Mb/s and the rate for low-priority traffic is set to 100 Mb/s, then when the customer transmits both high- and low-priority traffic, the aggregate amount of bandwidth consumed by customer traffic exceeds the contracted rate of 100 Mb/s and the transport provider transports more traffic than the contracted rate.

The second-tier shaper—that is, the per-SAP aggregate shaper—is used to limit the traffic at a configured rate on a per-SAP basis. The per-queue rates and behavior are not affected when the aggregate shaper is enabled. That is, as long as the aggregate rate is not reached then there are no changes to the behavior. If the aggregate rate limit is reached, then the per-SAP aggregate shaper throttles the traffic at the configured aggregate rate while preserving the 16-priority scheduling priorities that are used on shaped SAPs.

3.1.8 Per-VLAN network egress shapers

This section provides information about per-VLAN network egress shapers for Gen-2 adapter cards and platforms. For information about Gen-3 adapter cards and platforms, see [QoS for Gen-3 adapter cards and platforms](#).

The 7705 SAR supports a set of eight network egress queues on a per-port or on a per-VLAN basis for network Ethernet ports. Eight unique per-VLAN CoS queues are created for each VLAN when a per-VLAN shaper is enabled. When using per-VLAN shaper mode, in addition to the per-VLAN eight CoS queues, there is a single set of eight queues for hosting traffic from all unshaped VLANs, if any. VLAN shapers are enabled on a per-interface basis (that is, per VLAN) when a network queue policy is assigned to the interface. See [Per-VLAN shaper support](#) for a list of cards and nodes that support per-VLAN shapers.

On a network port with dot1q encapsulation, shaped and unshaped VLANs can coexist. In such a scenario, each shaped VLAN has its own set of eight CoS queues and is shaped with its own configured dual-rate shaper. The remaining VLANs (that is, the unshaped VLANs) are serviced using the **unshaped-if-cir** rate, which is configured using the **config>port>ethernet>network>egress>unshaped-if-cir** command. Assigning a rate to the unshaped VLANs is required for arbitration between the shaped VLANs and the bulk (aggregate) of unshaped VLANs, where each shaped VLAN has its own shaping rate while the aggregate of the unshaped VLANs has a single rate assigned to it.


Per-VLAN shapers are supported on dot1q-encapsulated ports. They are not supported on null- or qinq-encapsulated ports.

The following figure illustrates the queuing and scheduling blocks for network egress VLAN traffic.



Note: Because of space limitations in the figure, the second-tier, per-VLAN aggregate shapers are represented as a single loop containing the label “per VLAN”, even though they are dual-rate shapers similar to the third-tier network aggregate shaper.

Figure 5: Network egress shaped and unshaped VLAN queuing and scheduling



Note 1: Each per-VLAN aggregate shaper is a dual rate shaper similar to the 3rd tier dual rate shaper

24354

3.1.8.1 Shaped and unshaped VLANs

Shaped VLANs have user-configured rate limits (PIR and CIR)—called the aggregate rate limit—and must use 16-priority scheduling mode. Shaped VLANs operate on a per-interface basis and are enabled after a network queue policy is assigned to the interface. If a VLAN does not have a network queue policy assigned to the interface, it is considered an unshaped VLAN.

To configure a shaped VLAN with aggregate rate limits, use the **agg-rate-limit** command in the **config>router>if>egress** context. If the VLAN shaper is not enabled, the **agg-rate-limit** settings do not apply. The default aggregate rate limit (PIR) is set to the port egress rate.

Unshaped VLANs use default rate limits (PIR is the maximum possible port rate and CIR is 0 kb/s) and use 16-priority scheduling mode. All unshaped VLANs are classified, queued, buffered, and scheduled into an aggregate flow that gets prepared for third-tier arbitration by a single VLAN aggregate shaper.

In order for the aggregated unshaped VLANs to compete fairly for bandwidth with the shaped VLANs, a single, aggregate CIR for all the unshaped VLANs can be configured using the **unshaped-if-cir** command. The aggregate CIR is applied to all the unshaped VLANs as a group, not to individual VLANs, and is configured in the **config>port>ethernet> network>egress** context.

3.1.8.2 Per-VLAN shaper support

The following cards and nodes support network egress per-VLAN shapers:

- 6-port Ethernet 10Gbps Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card (1-port 10GigE mode and 10-port 1GigE mode)
- Packet Microwave Adapter card (includes 1+1 redundancy)
- v-port on the 2-port 10GigE (Ethernet) Adapter card/module
- 7705 SAR-A
- 7705 SAR-Ax
- 7705 SAR-H
- 7705 SAR-Hc
- 7705 SAR-M
- 7705 SAR-Wx
- 7705 SAR-X



Note: Per-VLAN network egress shapers are not supported on:

- Fast Ethernet ports (including ports 9 to 12 on the 7705 SAR-A)
- Gigabit Ethernet ports in Fast Ethernet mode
- non-datapath Ethernet ports (for example, the Management port)

3.1.8.3 VLAN shaper applications

This section describes the following two scenarios:

- [VLAN shapers for dual uplinks](#)
- [VLAN shapers for aggregation site](#)

3.1.8.3.1 VLAN shapers for dual uplinks

One of the main uses of per-VLAN network egress shapers is to enable load balancing across dual uplinks out of a spoke site. [Figure 6: VLAN shapers for dual uplinks](#) represents a typical hub-and-spoke mobile backhaul topology. To ensure high availability through the use of redundancy, a mobile operator invests in dual 7750 SR nodes at the MTSO. Dual 7750 SR nodes at the MTSO offer equipment protection, as well as protection against backhaul link failures.


In this example, the cell site 7705 SAR is dual-homed to 7750 SR_1 and SR_2 at the MTSO, using two disjoint Ethernet virtual connections (EVCs) leased from a transport provider. Typically, the EVCs have the

same capacity and operate in an forwarding/standby manner. One of the EVCs—the 7750 SR—transports all the user/mobile traffic to and from the cell site at any given time. The other EVC transports only minor volumes of control plane traffic between network elements (the 7705 SAR and the 7750 SR). Leasing two EVCs with the same capacity and using only one of them actively wastes bandwidth and is expensive (the mobile operator pays for two EVCs with the same capacity).

Mobile operators with increasing volumes of mobile traffic look for ways to use both of the EVCs simultaneously, in an active/active manner. In this case, using per-VLAN shapers would ensure that each EVC is loaded up to the leased capacity. Without per-VLAN shapers, the 7705 SAR supports a single per-port shaper, which does not meet the active/active requirement:

- If the egress rate is set to twice the EVC capacity, either one of the EVCs can end up with more traffic than its capacity.
- If the egress rate is set to the EVC capacity, half of the available bandwidth can never be consumed, which is similar to the 7705 SAR having no per-VLAN egress shapers.

Figure 6: VLAN shapers for dual uplinks




24355

3.1.8.3.2 VLAN shapers for aggregation site

Another typical use of per-VLAN shapers at network egress is shown in [Figure 7: VLAN shapers in aggregation site scenario](#). The figure shows a hub-and-spoke mobile backhaul network where EVCs leased from a transport provider are groomed to a single port, typically a 10-Gigabit Ethernet or a 1-Gigabit Ethernet port, at the hand-off point at the hub site. Traffic from different cell sites is handed off to the aggregation node over a single port, where each cell site is uniquely identified by the VLAN assigned to it.

In the network egress direction of the aggregation node, per-VLAN shaping is required to ensure traffic to different cell sites is shaped at the EVC rate. The EVC for each cell site would typically have a different rate. Therefore, every VLAN feeding into a particular EVC needs to be shaped at its own rate. For example, compare a relatively small cell site (Cell Site-1) at 20 Mb/s rate with a relatively large cell site (Cell Site-2) at 200 Mb/s rate. Without the granularity of per-VLAN shaping, shaping only at the per-port level cannot ensure that an individual EVC does not exceed its capacity.

Figure 7: VLAN shapers in aggregation site scenario



24356

3.1.9 Per-customer aggregate shapers (multiservice site) on Gen-2 hardware

This section provides information about per-customer aggregate shapers for Gen-2 adapter cards and platforms. For information about Gen-3 adapter cards and platforms, see [QoS for Gen-3 adapter cards and platforms](#).

A per-customer aggregate shaper is an aggregate shaper into which multiple SAP aggregate shapers can feed. The SAPs can be shaped at a desired rate called the Multiservice Site (MSS) aggregate rate. At ingress, SAPs that are bound to a per-customer aggregate shaper can span a whole Ethernet MDA meaning that SAPs mapped to the same MSS can reside on any port on an Ethernet MDA.

At egress, SAPs that are bound to a per-customer aggregate shaper can only span a port. Toward the fabric at ingress and toward the port at egress, multiple per-customer aggregate shapers are shaped at their respective configured rates to ensure fair sharing of available bandwidth among different per-customer aggregate shapers. Deep ingress queuing capability ensures that traffic bursts are absorbed, not dropped. Multi-tier shapers are based on an end-to-end backpressure mechanism that uses the following order (egress is given as an example):

- per-port egress rate (if configured), backpressures to
- per-customer aggregate shapers, backpressures to
- per-SAP aggregate shapers, backpressures to
- per-CoS queues (in the scheduling priority order)

To configure per-customer aggregate shaping, a shaper policy must be created and shaper groups must be created within that shaper policy. For access ingress per-customer aggregate shaping, a shaper policy must be assigned to an Ethernet MDA and SAPs on that Ethernet MDA must be bound to a shaper group within the shaper policy bound to that Ethernet MDA. For access egress per-customer aggregate shaping, a shaper policy must be assigned to a port and SAPs on that port must be bound to a shaper group within

the shaper policy bound to that port. The unshaped SAP shaper group within the policy provides the shaper rate for all the unshaped SAPs (4-priority scheduled SAPs). For each shaped SAP, however, an ingress or egress shaper group can be specified. For more information about shaper policies, see [Applying a shaper QoS policy and shaper groups](#).

The access ingress shaper policy is configured at the MDA level for fixed platforms. The default value for an access ingress shaper policy for each MDA and module is blank, as configured using the **no shaper-policy** command. On all 7705 SAR fixed platforms (with the exception of the 7705 SAR-X), when no MSS is configured, the existing access ingress aggregate rate is used as the shaper rate for the bulk of access ingress traffic. In order to use MSS, a shaper policy must be assigned to the access ingress interface of one MDA, and the shaper policy change is cascaded to all MDAs and modules in the chassis.

Before the access ingress shaper policy is assigned, the **config system qos access-ingress-aggregate-rate 10000000 unshaped-sap-cir max** command must be configured. Once a shaper policy is assigned to an access ingress MDA, the values configured using the **access-ingress-aggregate-rate** command cannot be changed.

On all 7705 SAR fixed platforms (with the exception of the 7705 SAR-X), when a shaper policy is assigned to an Ethernet MDA for access ingress aggregate shaping, it is automatically assigned to all the Ethernet MDAs in that chassis. The shaper group members contained in the shaper policy span all the Ethernet MDAs. SAPs on different Ethernet MDAs configured with the same ingress shaper group will share the shaper group rate.

Once the first MSS is configured, traffic from the designated SAPs is mapped to the MSS and shaped at the configured rate. The remainder of the traffic is then shaped according to the configured unshaped SAP rate. When a second MSS is added, SAPs that are mapped to the second MSS are shaped at the configured rate and traffic is arbitrated between the first MSS, the second MSS and unshaped SAP traffic.

In the access egress direction, the default shaper policy is assigned to each MSS-capable port. Ports that cannot support MSS are assigned a blank value, as configured using the **no shaper-policy** command. The default egress shaper group is assigned to each egress SAP that supports MSS. If the SAP does not support MSS, the egress SAP is assigned a blank value, as configured using the **no shaper-group** command.

3.1.9.1 MSS support

The following cards, modules, and platforms support MSS:

- 6-port Ethernet 10Gbps Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- Packet Microwave Adapter card
- 6-port SAR-M Ethernet module
- 7705 SAR-A
- 7705 SAR-Ax
- 7705 SAR-H
- 7705 SAR-Hc
- 7705 SAR-M
- 7705 SAR-Wx

- 7705 SAR-X



Note: MSS is not supported on the following:

- 4-port SAR-H Fast Ethernet module
- Fast Ethernet ports on the 7705 SAR-A

3.1.9.2 MSS and LAG interaction on the 7705 SAR-8 Shelf V2 and 7705 SAR-18

A SAP that uses a LAG can include two or more ports from the same adapter card or two different adapter cards.

In the access egress direction, each port can be assigned a shaper policy for access and can have shaper groups configured with different shaping rates. If a shaper group is not defined, the default shaper group is used. The port egress shaper policy, when configured on a LAG, must be configured on the primary LAG member. This shaper policy is propagated to each of the LAG port members, ensuring that each LAG port member has the same shaper policy.

The following egress MSS restrictions ensure that both active and standby LAG members have the same configuration:

- When a SAP is created using a LAG, whether the LAG has any port members and whether the egress scheduler mode is 4-priority or 16-priority, the default shaper group is automatically assigned to the SAP.
- Shaper groups cannot be changed from the default if they are assigned to SAPs that use LAGs with no port members.
- The last LAG port member cannot be deleted from a LAG that is used by any SAP that is assigned a non-default shaper group.
- The shaper policy or shaper group is not checked when the first port is added as a member of a LAG. When a second port is added as a member of a LAG, it can only be added if the shaper policy on the second port is the same as the shaper policy on the first member port of the LAG.
- The shaper group of a LAG SAP can be changed to a non-default shaper group only if the new shaper group exists in the shaper policy used by the active LAG port member.
- A shaper group cannot be deleted if it is assigned to unshaped SAPs (**unshaped-sap-shaper-group** command) or if it is used by at least one LAG SAP or non-LAG SAP.
- The shaper policy assigned to a port cannot be changed unless all of the SAPs on that port are assigned to the default shaper group.

In the ingress direction, there can be two different shaper policies on two different adapter cards for the two port members in a LAG. When assigning a shaper group to an ingress LAG SAP, each shaper policy assigned to the LAG port MDAs must contain that shaper group or the shaper group cannot be assigned. In addition, after a LAG activity switch occurs, the CIR/PIR configuration from the subgroup of the policy of the adapter card of the newly active member will be used.

The following ingress MSS restrictions allow the configuration of shaper groups for LAG SAPs, but the router ignores shaper groups that do not meet the restrictions:

- When a SAP is created using a LAG, whether the LAG has any port members and whether the ingress scheduler mode is 4-priority or 16-priority, the default shaper group is automatically assigned to the SAP.

- Shaper groups cannot be changed from the default if they are assigned to SAPs that use LAGs with no port members.
- The last LAG port member cannot be deleted from a LAG that is used by any SAP that is assigned a non-default shaper group.
- The shaper policy or shaper group is not checked when the first port is added as a member of a LAG. When a second port is added as a member of a LAG, all SAPs using the LAG are checked to ensure that any non-default shaper groups already configured on the SAPs are part of the shaper policy assigned to the adapter card of the second port. If the check fails, the port member is rejected.
- The shaper group of a LAG SAP can be changed to a non-default shaper group only if the new shaper group exists in the shaper policies used by all adapter cards of all LAG port members.
- A shaper group cannot be deleted if it is assigned to unshaped SAPs (**unshaped-sap-shaper-group** command) or if it is used by at least one LAG SAP or non-LAG SAP.
- The shaper policy assigned to an adapter card cannot be changed unless all of the SAPs on that adapter card are assigned to the default shaper group.

3.1.10 QoS for hybrid ports on Gen-2 hardware

This section provides information about QoS for hybrid ports on Gen-2 adapter cards and platforms. For information about Gen-3 adapter cards and platforms, see [QoS for Gen-3 adapter cards and platforms](#).

In the ingress direction of a hybrid port, traffic management behavior is the same as it is for access and network ports. See [Access ingress](#) and [Network ingress](#).

In the egress direction of a hybrid port, access and network aggregate shapers are used to arbitrate between the bulk (aggregate) of access and network traffic flows. As shown in [Figure 8: Hybrid port egress shapers and schedulers on Gen-2 hardware](#), on the access side (above the solid line), both the access egress SAP aggregates (#1) and the unshaped SAP shaper (#2) feed into the access egress aggregate shaper (#3). On the network side (below the solid line), both the per-VLAN shapers (#4) and the unshaped interface shaper (#5) feed into the network egress aggregate shaper (#6). Then, the access and the network aggregate shapers are arbitrated in a dual-rate manner, in accordance with their respective configured committed and peak rates (#7). As a last step, the **egress-rate** for the port (when configured) applies backpressure to both the access and the network aggregate shapers, which apply backpressure all the way to the FC queues belonging to both access and network traffic.



Note: Because of space limitations in the figure, the second-tier, per-SAP and per-VLAN aggregate shapers are represented as a single loop containing the label "per SAP" or "per VLAN", even though they are dual-rate shapers similar to the third-tier network aggregate shaper. Tiers are labeled at the top of the figure.

As part of the hybrid port traffic management solution, access and network second-tier shapers are bound to access and network aggregate shapers, respectively. The hybrid port egress datapath can be visualized as access and network datapaths that coexist separately up until the access and network aggregate shapers at Tier 3 (#3 and #6).

In the figure, the top half is identical to access egress traffic management, where CoS queues (Tier 1) feed into either per-SAP shapers for shaped SAPs (#1) or a single second-tier shaper for all unshaped SAPs (#2). Up to the end of the second-tier, per-SAP aggregate shapers, the access egress datapath is maintained in the same manner as an Ethernet port in access mode. The same logic applies for network egress. The bottom half of the figure shows the datapath from the CoS queues to the per-VLAN shapers, which is identical to the datapath for any other Ethernet port in network mode.

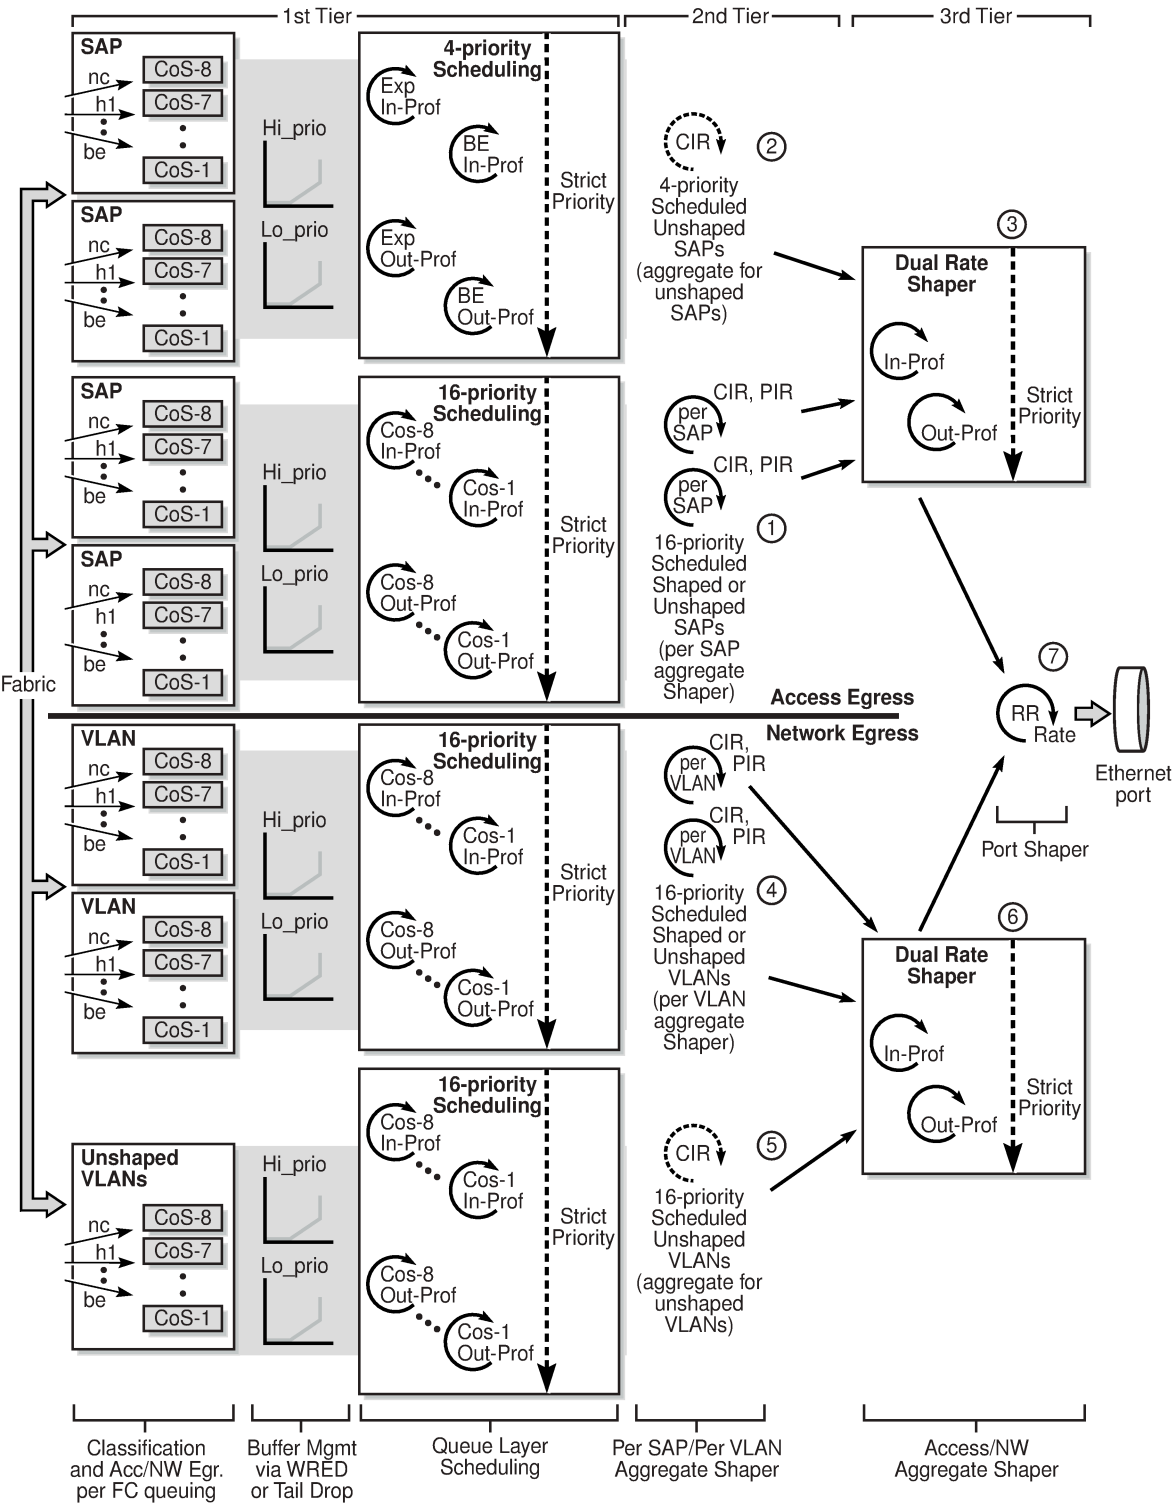
The main difference between hybrid mode and access and network modes is shown when the access and the network traffic is arbitrated toward the port (Tier 3). At this point, a new set of dual-rate shapers (called shaper groups) are introduced: one shaper for the aggregate (bulk) of the access traffic (#3) and another shaper for the aggregate of the network traffic (#6), to ensure rate-based arbitration among access and network traffic.

Depending on the use and the application, the committed rate for any one mode of flow may need to be fine-tuned to minimize delay, jitter and loss. In addition, through the use of egress-rate limiting, a fourth level of shaping can be achieved.

When **egress-rate** is configured (under **config>port>ethernet**), the following events occur:

- **egress-rate** applies backpressure to the access and network aggregate shapers
- as a result, the aggregate shapers apply backpressure to the per-SAP and per-VLAN aggregate shapers
 - access aggregate shapers apply backpressure to the per-SAP aggregate shapers and the unshaped SAP aggregate shaper
 - network aggregate shapers apply backpressure to the per-VLAN aggregate shapers and the unshaped VLAN aggregate shaper
- as a result, the per-SAP and per-VLAN aggregate shapers apply backpressure to their respective CoS queues

Figure 8: Hybrid port egress shapers and schedulers on Gen-2 hardware



24073

3.1.11 QoS for Gen-3 adapter cards and platforms

Third-generation (Gen-3) Ethernet adapter cards and Ethernet ports on Gen-3 platforms support 4-priority scheduling.

The main differences between Gen-3 hardware and Gen-2 hardware are that on Gen-3 hardware:

- SAPs are shaped (that is, no unshaped SAPs)
- SAPs and VLANs are shaped by 4-priority schedulers, not 16-priority schedulers
- 4-priority scheduling is done on a per-SAP basis
- backpressure is applied according to relative priority across VLANs and interfaces. That is, scheduling is carried out in priority order, ignoring per-VLAN and per-interface boundaries. Conforming, expedited traffic across all queues is scheduled regardless of the VLAN boundaries. After all the conforming, expedited traffic across all queues has been serviced, the servicing of conforming, best effort traffic begins.

See [Scheduling modes](#) for a summary of scheduler mode support. For information about adapter card generations, see the “Evolution of Ethernet Adapter Cards, Modules, and Platforms” section in the 7705 SAR Interface Configuration Guide.

The following figure describes the access, network, and hybrid port scheduling behavior for Gen-3 hardware and compares it with the scheduling behavior of Gen-2 hardware.

Table 5: Scheduling behavior for Gen-2 and Gen-3 hardware

Port type		Gen-3 hardware with 4-priority mode	Gen-2 hardware with 4-priority mode
Access	Within a SAP	EXP over BE	EXP over BE
	Default configuration	Simple round-robin (RR) scheduling among SAPs	EXP (across all queues, no SAP boundaries) over BE
	H-QoS and MSS aggregate shapers	RR among aggregates based on PIR and CIR (SAP at tier 2, MSS at tier 3)	N/A
Network	Default configuration (8 queues per port)	EXP over BE	EXP over BE
	Per-VLAN shaper	EXP over BE	RR among VLAN shapers based on PIR and CIR
Hybrid	Default configuration (8 queues per port)	EXP over BE	EXP (across all SAPs and network queues) over BE
	Per-VLAN shaper	RR among VLAN shapers based on PIR and CIR	RR among VLAN shapers based on PIR and CIR

In summary, the following updates to Gen-3 scheduling are implemented:

- enabled CIR-based shaping for:
 - per-SAP aggregate ingress and egress shapers

- per-customer aggregate ingress and egress shapers
- per-VLAN shaper at network egress when port is in hybrid mode
- access and network aggregate shapers for hybrid ports
- disabled backpressure to the FC queues dependent on the relative priority among all VLAN-bound IP Interfaces at:
 - access ingress and access egress when port is in access mode
 - access ingress and access egress when port is in hybrid mode
 - network ingress
 - network egress when port is in hybrid mode



Note: For network egress traffic when the port is in network mode, there is no change to CIR-based shaping for per-VLAN shapers (that is, PIR-based shaping only, CIR-based shaping is not enabled), and backpressure to the FC queues based on the relative priority among all VLAN-bound IP interfaces is still enabled.

3.1.11.1 6-port SAR-M Ethernet module

The egress datapath shapers on a 6-port SAR-M Ethernet module operate on the same frame size as any other shaper. These egress datapath shapers are:

- per-queue shapers
- per-SAP aggregate shapers
- per-customer aggregate (MSS) shapers

The egress port shaper on a 6-port SAR-M Ethernet module does not account for the 4-byte FCS. Packet byte offset can be used to make adjustments to match the desired operational rate or eliminate the implications of FCS. See [Packet byte offset](#) for more information.

3.1.11.2 4-priority scheduling behavior on Gen-3 hardware

For access ingress, access egress, and network ingress traffic, the behavior of 4-priority scheduling on Gen-3 hardware is similar to 4-priority scheduling on Gen-2 hardware. See the following figures:

- [Figure 9: 4-priority scheduling at access ingress \(Gen-3 hardware\)](#) (access ingress)
- [Figure 10: 4-priority scheduling at access egress \(Gen-3 hardware\)](#) (access egress)
- [Figure 11: 4-priority scheduling at network ingress \(Gen-3 hardware\): per-destination mode](#) (network ingress, destination mode)
- [Figure 12: 4-priority scheduling at network ingress \(Gen-3 hardware\): aggregate mode](#) (network ingress, aggregate mode)

For network egress traffic through a network port on Gen-3 hardware, the behavior of 4-priority scheduling is as follows: traffic priority is determined at the queue-level scheduler, which is based on the queue PIR and CIR and the queue type. The queue-level priority is carried through the various shaping stages and is used by the 4-priority Gen-3 VLAN scheduler at network egress. See [Figure 13: 4-priority scheduling at network egress \(Gen-3 hardware\) on a network port](#) and its accompanying description.

For hybrid ports, both access and network egress traffic use 4-priority scheduling that is similar to 4-priority scheduling on Gen-2 hardware. See [Figure 14: 4-priority scheduling for hybrid port egress \(Gen-3 hardware\)](#) and its accompanying description.

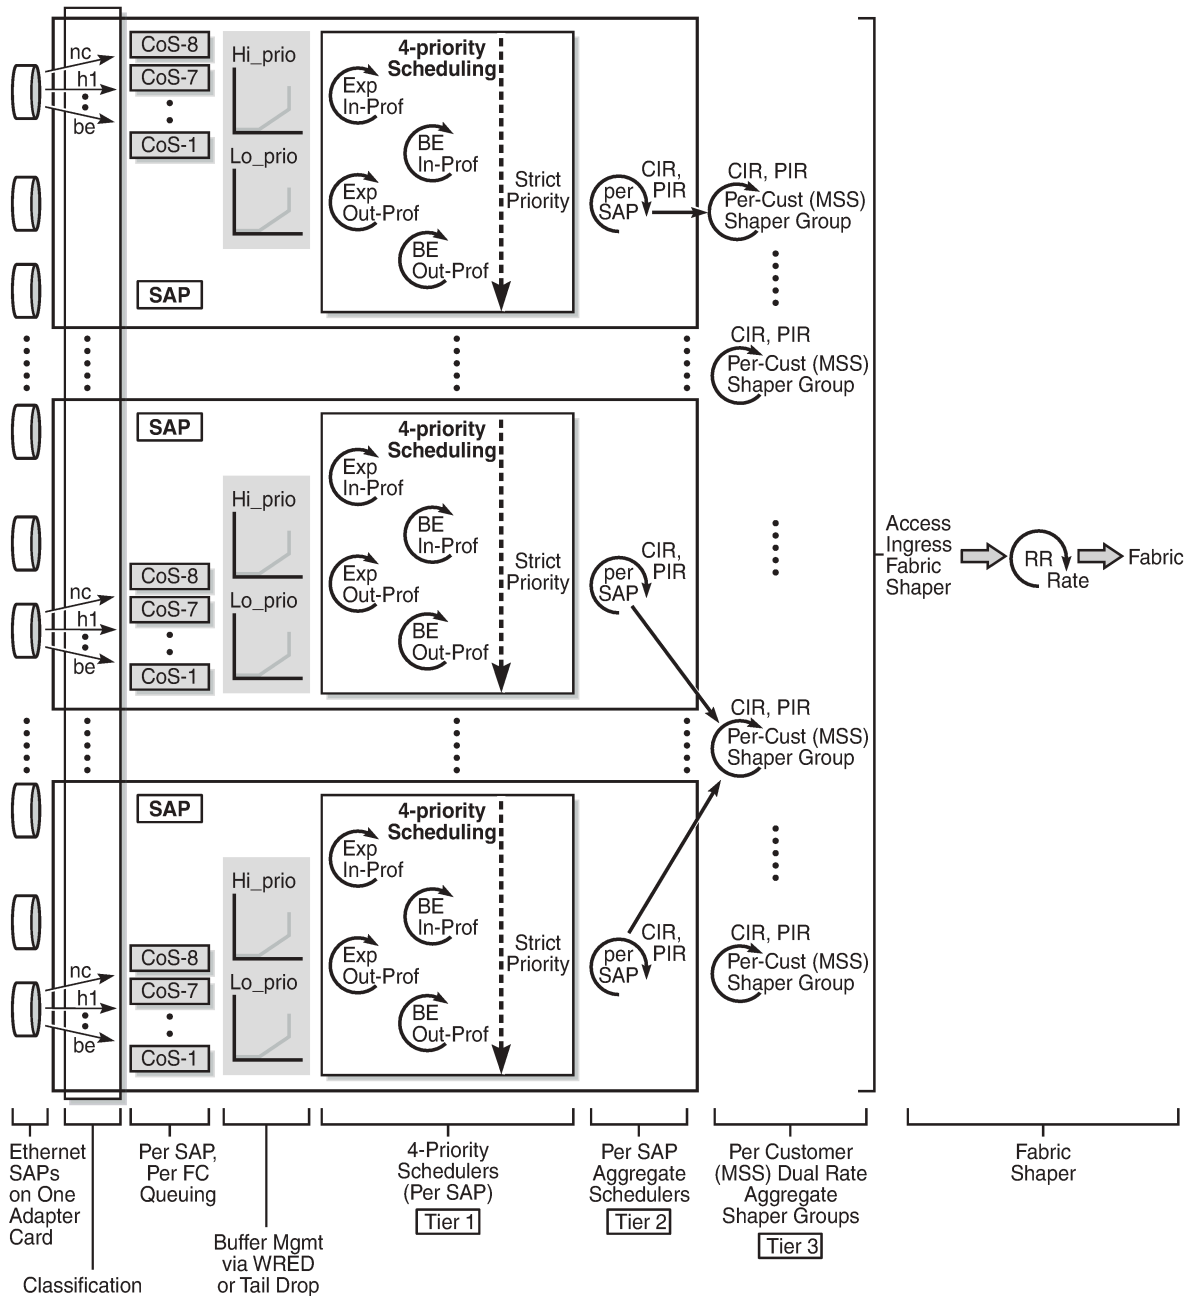


Note: The 7705 SAR-X defaults to 4-priority mode and does not support fabric shapers. Traffic from fabric shapers (access and network) are arbitrated in a round-robin manner toward the egress datapath.

In the following figure, the shaper groups all belong within one shaper policy and only one shaper policy is assigned to an ingress adapter card. Each SAP can be associated with one shaper group. Multiple SAPs can be associated with the same shaper group. All the access ingress traffic flows to the access ingress fabric shaper, in-profile (conforming) traffic first, then out-of-profile (non-conforming) traffic. Network ingress traffic functions similarly.

The 4-priority schedulers on Gen-2 and Gen-3 hardware are very similar, except that 4-priority scheduling on Gen-3 hardware is done on a per-SAP basis.

Figure 9: 4-priority scheduling at access ingress (Gen-3 hardware)

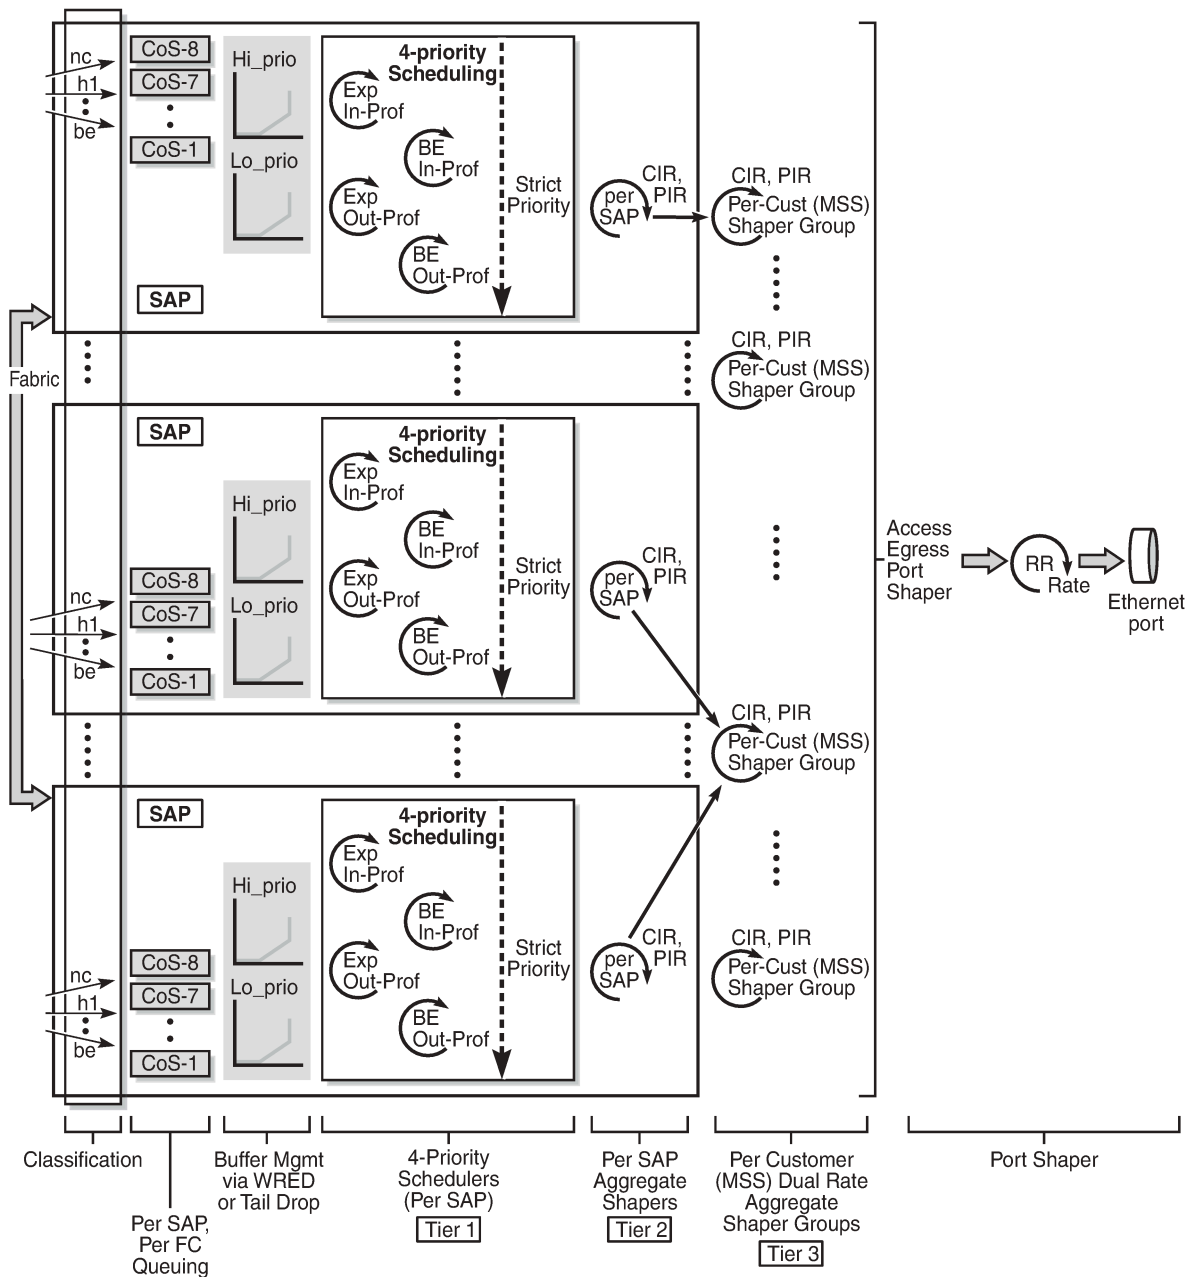


25887

The following figure shows 4-priority scheduling for access egress on Gen-3 hardware. QoS behavior for access egress is similar to QoS behavior for access ingress.

The 4-priority schedulers on Gen-2 and Gen-3 hardware are very similar, except that 4-priority scheduling on Gen-3 hardware is done on a per-SAP basis.

Figure 10: 4-priority scheduling at access egress (Gen-3 hardware)

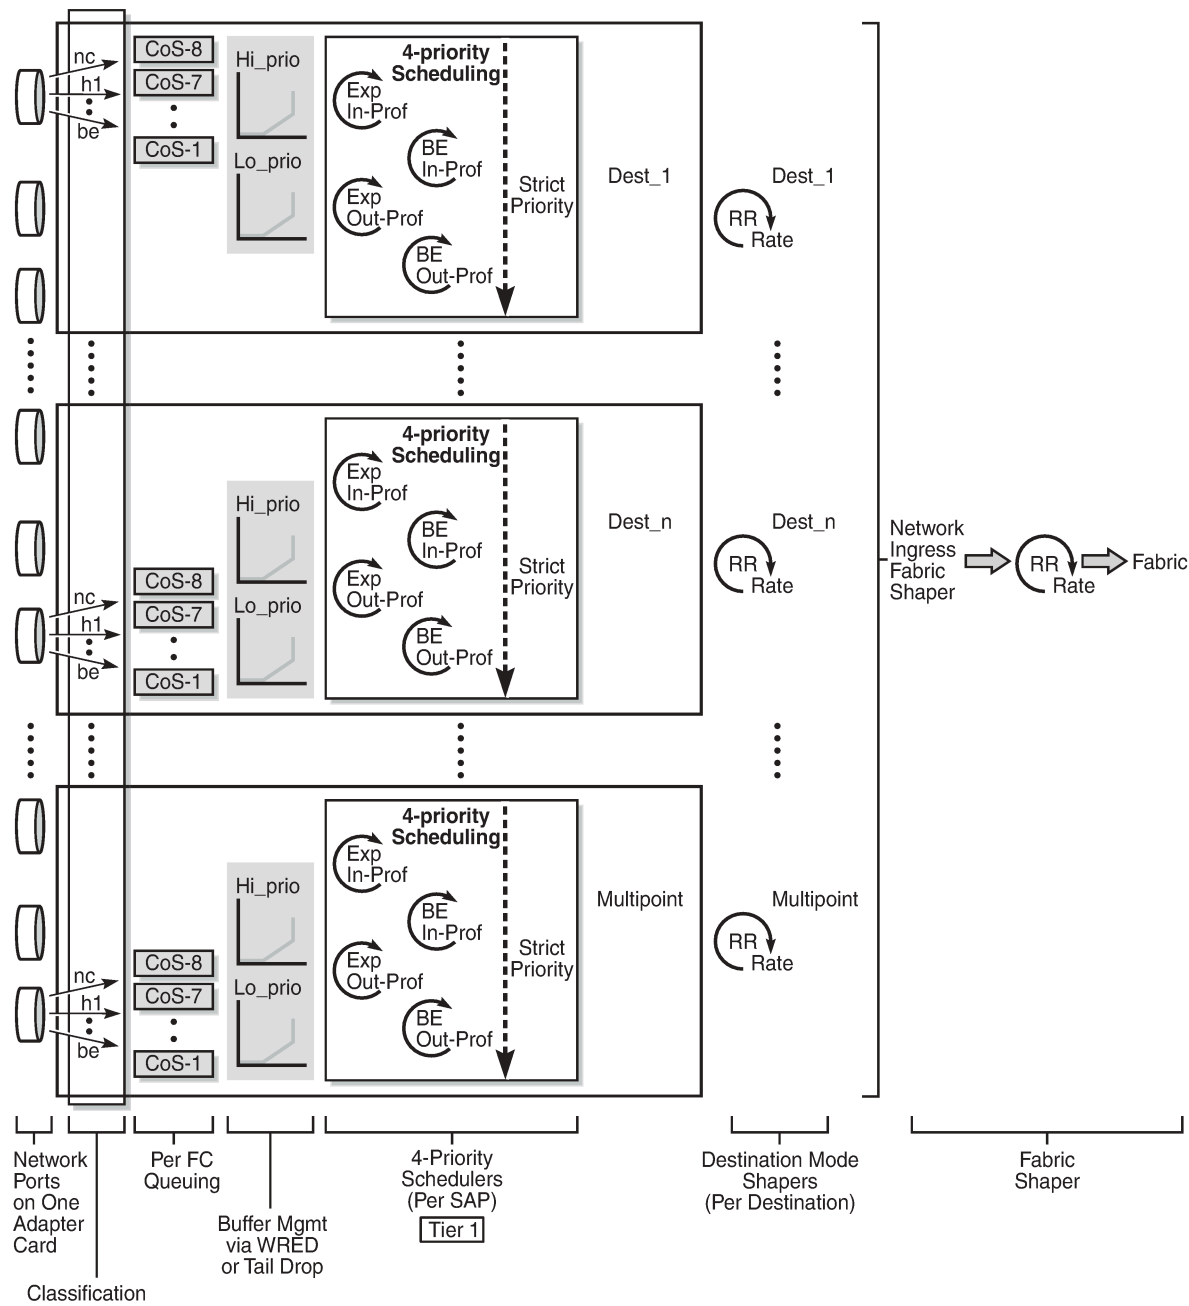


25888

Figure 11: 4-priority scheduling at network ingress (Gen-3 hardware): per-destination mode and Figure 12: 4-priority scheduling at network ingress (Gen-3 hardware): aggregate mode show network ingress scheduling for per-destination and aggregate modes, which are configured under the **fabric-profile** command. Traffic arriving on a network port is examined for its destination MDA and directed to the QoS block that sends traffic to the appropriate MDA. There is one set of queues for each block, and an additional set for multipoint traffic.

In the following figure, there is one per-destination shaper for each destination MDA.

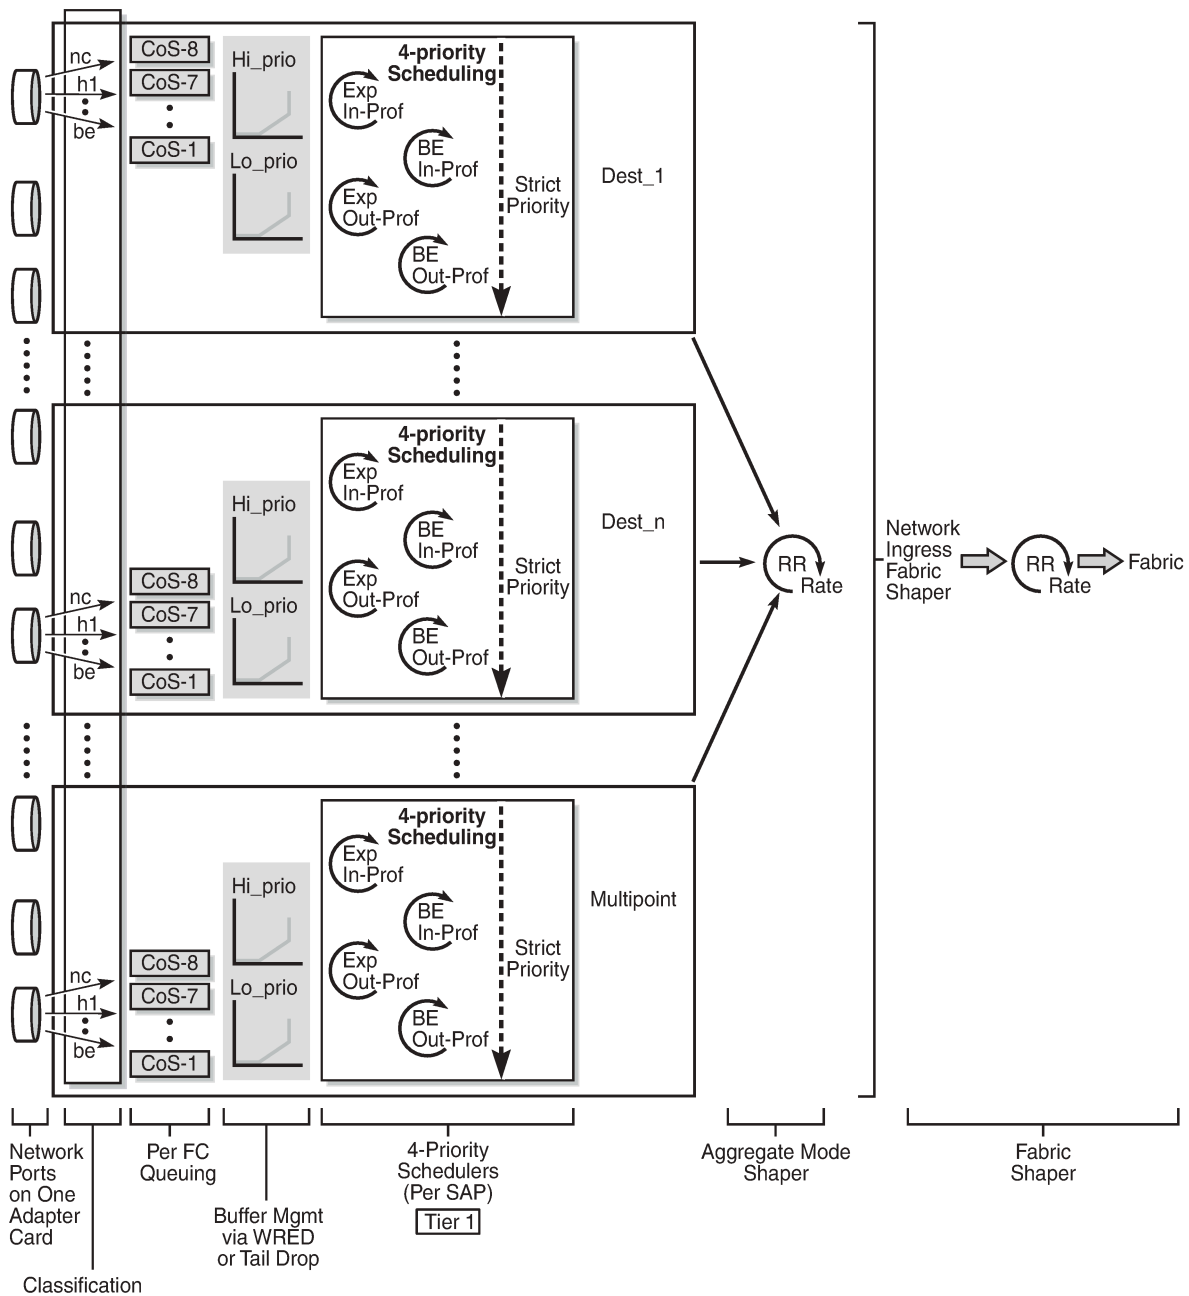
Figure 11: 4-priority scheduling at network ingress (Gen-3 hardware): per-destination mode



25889

In the following figure, there is a single shaper to handle all the traffic.

Figure 12: 4-priority scheduling at network ingress (Gen-3 hardware): aggregate mode



25890

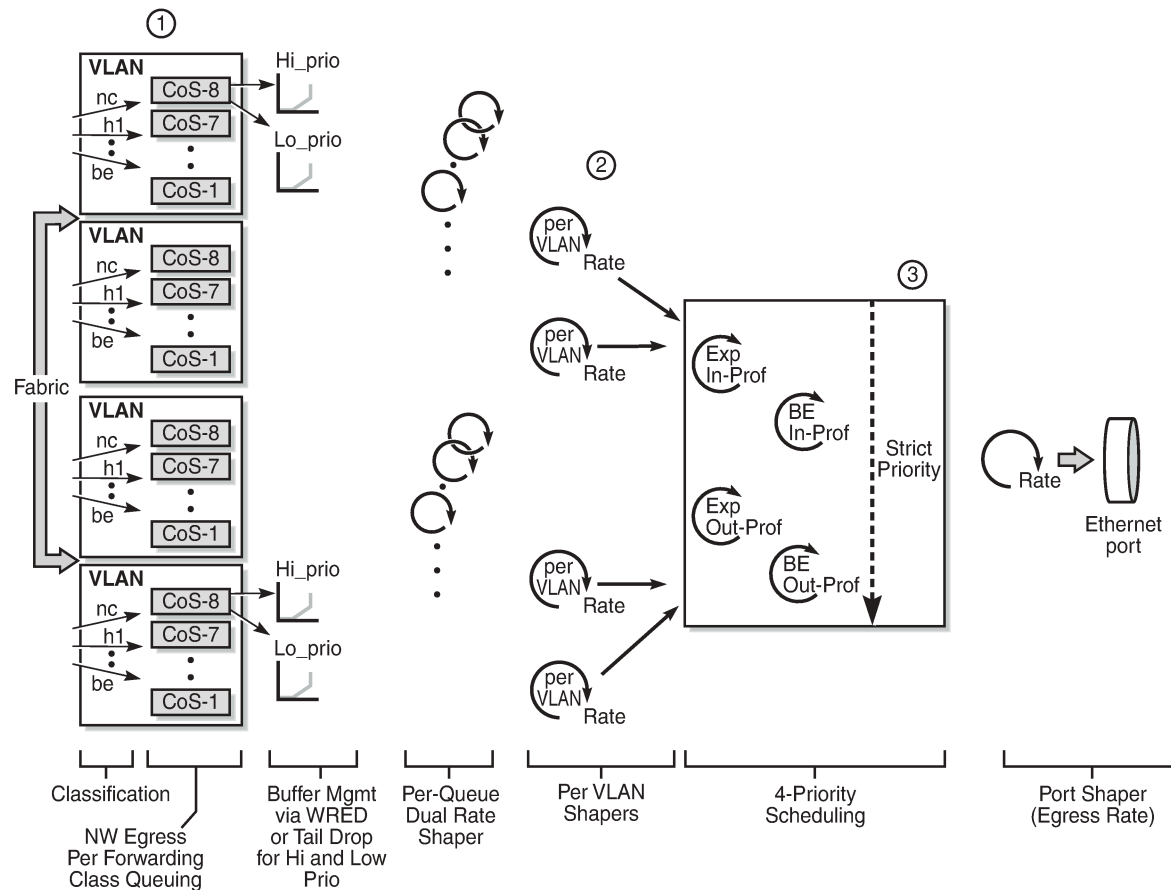
The following figure shows 4-priority scheduling for Gen-3 hardware at network egress. Queue-level CIR and PIR values and the queue type are determined at queuing and provide the scheduling priority for a specific flow across all shapers toward the egress port (#1 in the figure). At the per-VLAN aggregate level (#2), only a single rate—the total aggregate rate (PIR)—can be configured; CIR configuration is not supported at the per-VLAN aggregate-shaper level for network egress traffic. All VLANs are aggregated and scheduled by a 4-priority aggregate scheduler (#3). The flow is then fed to the port shaper and

processed at the egress rate. In case of congestion, the port shaper provides backpressure, resulting in the buffering of traffic by individual FC queues until the congested state ends.



Note: The behavior of 4-priority scheduling for network egress traffic through a network port differs from the behavior for network egress traffic through a hybrid port. The following figures show the differences.

Figure 13: 4-priority scheduling at network egress (Gen-3 hardware) on a network port

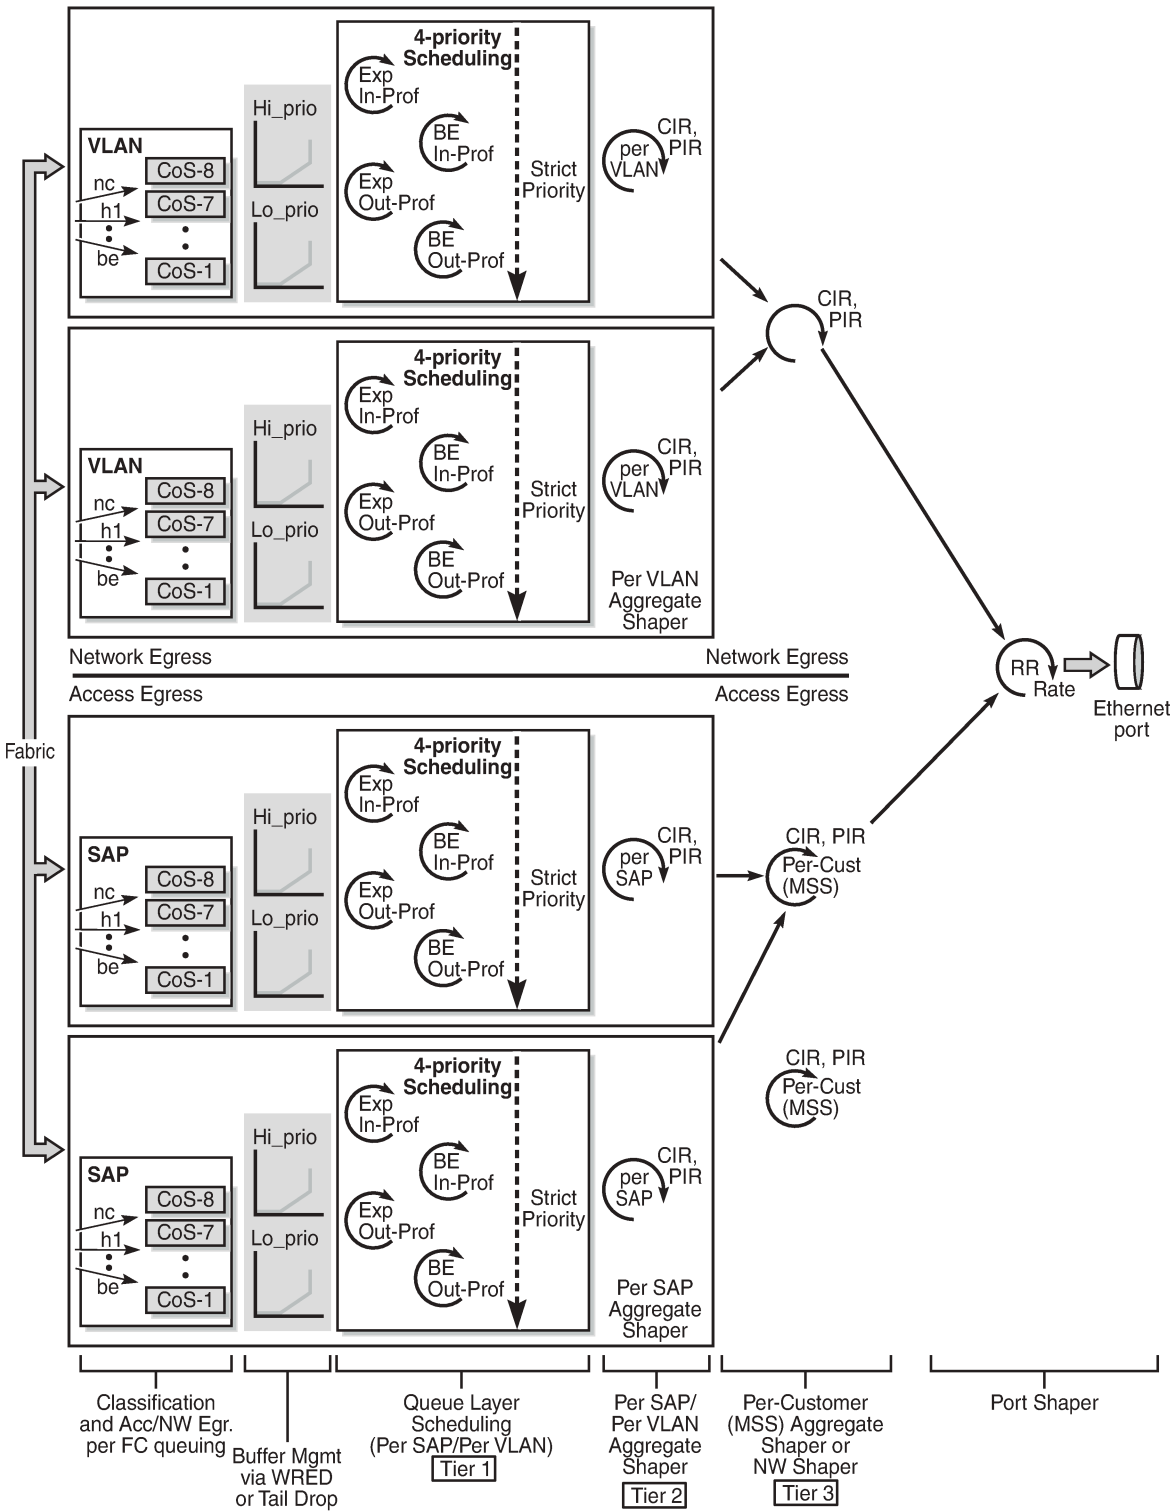


24775

The following figure shows 4-priority scheduling for Gen-3 hardware where ports are in hybrid mode. The QoS behavior for both access and network egress traffic is similar except that the access egress path includes tier 3, per-customer aggregate shapers. Access and network shapers prepare and present traffic to the port shaper, which arbitrates between access and network traffic.

The 4-priority schedulers on the Gen-2 and Gen-3 hardware are very similar, except that 4-priority scheduling on Gen-3 hardware is done on a per-SAP or a per-VLAN basis (for access egress and network egress, respectively).

Figure 14: 4-priority scheduling for hybrid port egress (Gen-3 hardware)



25891

3.1.11.3 Gen-3 hardware and LAG

When a Gen-3-based port and a Gen-2-based port are attached to a LAG SAP (also referred to as mix-and-match LAG), configuring scheduler mode for the LAG SAP is required because it is used by Gen-2 ports, but it is ignored by the Gen-3 port.

For more information, see the "LAG Support on Third-Generation Ethernet Adapter Cards, Ports, and Platforms" section in the 7705 SAR Interface Configuration Guide.

3.1.12 QoS on a ring adapter card or module

This section contains overview information as well as information about the following topics:

- [Network and network queue QoS policy types](#)
- [Network QoS and network queue policies on a ring adapter card or module](#)
- [Considerations for using ring adapter card or module QoS policies](#)


The following figure shows a simplified diagram of the ports on a 2-port 10GigE (Ethernet) Adapter card (also known as a ring adapter card). The ports can also be conceptualized the same way for a 2-port 10GigE (Ethernet) module. A ring adapter card or module has physical Ethernet ports used for Ethernet bridging in a ring network (labeled Port 1 and Port 2 in the figure). These ports are referred to as the ring ports because they connect to the Ethernet ring. The ring ports operate on the Layer 2 bridging domain side of the ring adapter card or module, as does the add/drop port, which is an internal port on the card or module.

On the Layer 3 IP domain side of a ring adapter card or module, there is a virtual port (v-port) and a fabric port. The v-port is also an internal port. Its function is to help control traffic on the IP domain side of a ring adapter card or module.

To manage ring and add/drop traffic mapping to queues in the Layer 2 bridging domain, a ring type network QoS policy can be configured for the ring at the adapter card level (under the **config>card>mda** context). To manage ring and add/drop traffic queuing and scheduling in the Layer 2 bridging domain, network queue QoS policies can be configured for the ring ports and the add-drop port.

To manage add/drop traffic classification and remarking in the Layer 3 IP domain, IP-interface type network QoS policies can be configured for router interfaces on the v-port. To manage add/drop traffic queuing and scheduling in the Layer 3 IP domain, network queue QoS policies can be configured for the v-port and at network ingress at the adapter card level (under the **config>card>mda** context).

Figure 15: Ports on a 2-port 10GigE (Ethernet) Adapter card



23478

3.1.12.1 Network and network queue QoS policy types

All ports on a ring adapter card or module are possible congestion points and therefore can have network queue QoS policies applied to them.

In the bridging domain, a single ring type network QoS policy can be applied at the adapter card level and operates on the ring ports and the add/drop port. In the IP domain, IP interface type network QoS policies can be applied to router interfaces.

Network QoS policies are created using the **config>qos>network** command, which includes the **network-policy-type** keyword to specify the type of policy:

- ring (for bridging domain policies)
- ip-interface (for network ingress and network egress IP domain policies)
- default (for setting the policy to policy 1, the system default policy)

When the policy has been created, its default action and classification mapping can be configured.

3.1.12.2 Network QoS and network queue policies on a ring adapter card or module

Network QoS policies are applied to the ring ports and the add/drop port using the **qos-policy** command found under the **config>card>mda** context. These ports are not explicitly specified in the command.

Network queue QoS policies are applied to the ring ports and the v-port using the **queue-policy** command found under the **config>port** context. Similarly, a network queue policy is applied to the add/drop port using the **add-drop-port-queue-policy** command, found under the **config>card>mda** context. The add/drop port is not explicitly specified in this command.

The CLI commands for applying QoS policies are listed in this guide. The CLI command descriptions are in the 7705 SAR Interface Configuration Guide.

3.1.12.3 Considerations for using ring adapter card or module QoS policies

The following notes apply to configuring and applying QoS policies to a ring adapter card or module, as well as other adapter cards:

1. The ring ports and the add/drop port cannot use a network queue policy that is being used by the v-port or the network ingress port or any other port on other cards, and vice versa. This does not apply to the default network queue policy.
2. If a network-queue policy is assigned to a ring port or the add/drop port, all queues that are configured in the network-queue policy are created regardless of any FC mapping to the queues. All FC queue mapping in the network-queue policy is meaningless and is ignored.
3. If the QoS policy has a dot1p value mapped to a queue that is not configured in the network-queue policy, the traffic of this dot1p value is sent to queue 1.
4. If a dot1q-to-queue mapping is defined in a network QoS policy, and if the queue is not configured on any ring port or the add/drop port, all traffic received from a port will be sent out from queue 1 on the other two ports. For example, if traffic is received on port 1, it will be sent out on port 2 and/or the add/drop port.
5. Upon provisioning an MDA slot for a ring adapter card or module (**config>card>mda> mda-type**) an additional eight network ingress queues are allocated to account for the network queues needed for the add/drop port. When the ring adapter card or module is deprovisioned, the eight queues are deallocated.
6. The number of ingress network queues is shown by using the **tools>dump> system-resource** command.

3.1.13 QoS for IPSec traffic

For specific information about QoS for IPSec traffic, see the "QoS" section in the "IPSec" chapter in the 7705 SAR Services Guide.

3.1.14 QoS for network group encryption traffic

The 7705 SAR provides priority and scheduling for traffic into the encryption and decryption engines on nodes that support network group encryption (NGE). This applies to traffic at network ingress or network egress.

For specific information, see the "QoS for NGE Traffic" section in the "Network Group Encryption" chapter in the 7705 SAR Services Guide.

3.2 Access ingress

This section contains the following topics for traffic flow in the access ingress direction:

- [Access ingress traffic classification](#)
- [Access ingress queues](#)
- [Access ingress queuing and scheduling](#)

- [Access ingress per-SAP aggregate shapers \(access ingress H-QoS\)](#)
- [Ingress shaping to fabric \(access and network\)](#)
- [Configurable ingress shaping to fabric \(access and network\)](#)
- [Fabric shaping on the fixed platforms \(access and network\)](#)

3.2.1 Access ingress traffic classification

Traffic classification identifies a traffic flow and maps the packets belonging to that flow to a preset forwarding class, so that the flow can receive the required special treatment. Up to eight forwarding classes are supported for traffic classification. See [Table 2: Default forwarding classes](#) for a list of these forwarding classes.

For TDM channel groups, all of the traffic is mapped to a single forwarding class. Similarly, for ATM VCs, each VC is linked to one forwarding class. On Ethernet ports and VLANs, up to eight forwarding classes can be configured based on 802.1p (dot1p) bits or DSCP bits classification. On PPP/MLPPP, FR (for Lpipes), or cHDLc SAPs, up to eight forwarding classes can be configured based on DSCP bits classification. FR (for Fpipes) and HDLC SAPs are mapped to one forwarding class.



Note:

- If an Ethernet port is set to null encapsulation, the dot1p value has no meaning and cannot be used for classification purposes.
- If a port or SAP is set to qinq encapsulation, use the **match-qinq-dot1p top | bottom** command to indicate which qtag contains the dot1p bits that are used for classification purposes. The **match-qinq-dot1p** command is found under the **config>service** context. See the 7705 SAR Services Guide, "VLL Services Command Reference", for details.

After the classification takes place, forwarding classes are mapped to queues as described in the sections that follow.

3.2.1.1 Traffic classification types

The various traffic classification methods used on the 7705 SAR are described in the following table. A list of classification rules follows the table.

Table 6: Traffic classification types

Traffic classification based on...	Description
a channel group (n × DS0)	Applies to 16-port T1/E1 ASAP Adapter card and 32-port T1/E1 ASAP Adapter card ports, 2-port OC3/STM1 Channelized Adapter card ports, 12-port Serial Data Interface card ports, 4-port T1/E1 and RS-232 Combination module ports, and 6-port E&M Adapter card ports in structured or unstructured circuit emulation mode. In this mode, a number of DS0s are transported within the payload of the same Circuit Emulation over Packet Switched Networks (CESoPSN) packet, Circuit Emulation over Ethernet (CESoETH) packet, or Structure-Agnostic TDM over Packet (SAToP) packet. Therefore, the timeslots transporting the same type of traffic are classified simultaneously.

Traffic classification based on...	Description
an ATM VCI	<p>On ATM-configured ports, any virtual connection regardless of service category is mapped to the configured forwarding class. One-to-one mapping is the only supported option.</p> <p>VP- or VC-based classifications are both supported. A VC with a specified VPI and VCI is mapped to the configured forwarding class. A VP connection with a specified VPI is mapped to the configured forwarding class.</p>
an ATM service category	<p>Similar ATM service categories can be mapped against the same forwarding class. Traffic from a VC with a specified service category is mapped to the configured forwarding class. VC selection is based on the ATM VC identifier.</p>
an Ethernet port	<p>All the traffic from an access ingress Ethernet port is mapped to the selected forwarding class. More granular classification can be performed based on dot1p or DSCP bits of the incoming packets. Classification rules applied to traffic flows on Ethernet ports function in the same way as access/filter lists. There can be multiple tiers of classification rules associated with an Ethernet port. In this case, classification is performed based on priority of classifier. The order of the priorities is described in Hierarchy of classification rules.</p>
an Ethernet VLAN (dot1q or qinq)	<p>Traffic from an access Ethernet VLAN (dot1q or qinq) interface can be mapped to a forwarding class. Each VLAN can be mapped to one forwarding class.</p>
IEEE 802.1p bits (dot1p)	<p>The dot1p bits in the Ethernet/VLAN ingress packet headers are used to map the traffic to up to eight forwarding classes.</p>
PPP/MLPPP, FR (for Ipipes), and cHDLC SAPs	<p>Traffic from an access ingress SAP is mapped to the selected forwarding class. More granular classification can be performed based on DSCP bits of the incoming packets.</p>
FR (for Fpipes) and HDLC SAPs	<p>Traffic from an access ingress SAP is mapped to the selected (default) forwarding class.</p>
DSCP bits	<p>When the Ethernet payload is IP, ingress traffic can be mapped to a maximum of eight forwarding classes based on DSCP bit values.</p> <p>DSCP-based classification supports untagged, single-tagged, double-tagged, and triple-tagged Ethernet frames. If an ingress frame has more than three VLAN tags, then dot1q or qinq dot1p-based classification must be used.</p>
Multi-field classifiers	<p>Traffic is classified based on any IP criteria currently supported by the 7705 SAR filter policies; for example, source and destination IP address, source and destination port, if the packet is fragmented, ICMP code, and TCP state. For information about multi-field classification, see the 7705 SAR Router Configuration Guide, "Multi-field Classification (MFC)" and "IP, MAC, and VLAN Filter Entry Commands".</p>

3.2.1.1.1 Hierarchy of classification rules

The following table shows classification options for various access entities (SAP identifiers) and service types. For example, traffic from a TDM port using a TDM (Cpipe) PW maps to one FC (all traffic has the same CoS). Traffic from an Ethernet port using a Epipe PW can be classified to as many as eight FCs based on DSCP classification rules, while traffic from a SAP with dot1q or qinq encapsulation can be classified to up to eight FCs based on dot1p or DSCP rules.

For Ethernet traffic, dot1p-based classification for dot1q or QinQ SAPs takes precedence over DSCP-based classification. For null-encapsulated Ethernet ports, only DSCP-based classification applies. In either case, when defining classification rules, a more specific match rule is always preferred to a general match rule.

For more information about hierarchy rules, see [Table 19: Forwarding class and enqueueing priority classification hierarchy based on rule type](#) in the [Service ingress QoS policies](#) section.

Table 7: Access ingress traffic classification for SAPs per service type

Access type (SAP)	Service type							
	TDM PW	ATM PW	FR PW	HDLC PW	Ethernet PW	IP PW	VPLS	VPRN
TDM port	1 FC							
Channel group	1 FC							
ATM virtual connection identifier		1 FC					1 FC	
FR			1 FC			DSCP, up to 8 FCs		
HDLC				1 FC				
PPP / MLPPP						DSCP, up to 8 FCs		DSCP, up to 8 FCs
cHDLC						DSCP, up to 8 FCs		
Ethernet port					DSCP, up to 8 FCs	DSCP, up to 8 FCs	DSCP, up to 8 FCs	DSCP, up to 8 FCs
Dot1q encapsulation					Dot1p or DSCP, up to 8 FCs	Dot1p or DSCP, up to 8 FCs	Dot1p or DSCP, up to 8 FCs	Dot1p or DSCP, up to 8 FCs
Qinq encapsulation					Dot1p or DSCP, up to 8 FCs	Dot1p or DSCP, up to 8 FCs	Dot1p or DSCP, up to 8 FCs	Dot1p or DSCP, up to 8 FCs

3.2.1.1.2 Discard probability of classified traffic

When the traffic is mapped against a forwarding class, the discard probability for the traffic can be configured as high or low priority at ingress. When the traffic is further classified as high or low priority, different congestion management schemes could be applied based on this priority. For example WRED curves can then be run against the high- and low-priority traffic separately, as described in [Slope policies \(WRED and RED\)](#).

This ability to specify the discard probability is very significant because it controls the amount of traffic that is discarded under congestion or high usage. If you know the characteristics of your traffic, particularly the burst characteristics, the ability to change the discard probability can be used to great advantage. The objective is to customize the properties of the random discard functionality such that the minimal amount of data is discarded.

3.2.2 Access ingress queues

After the traffic is classified to different forwarding classes, the next step is to create the ingress queues and bind forwarding classes to these queues.

There is no restriction of a one-to-one association between a forwarding class and a queue. That is, more than one forwarding class can be mapped to the same queue. This capability is beneficial in that it allows a bulk-sum amount of resources to be allocated to traffic flows of a similar nature. For example, in the case of 3G UMTS services, HSDPA and OAM traffic are both considered BE in nature. However, HSDPA traffic can be mapped to a better forwarding class (such as L2) while OAM traffic can remain mapped to a BE forwarding class. But they both can be mapped to a single queue to control the total amount of resources for the aggregate of the two flows.

A large but finite amount of memory is available for the queues. Within this memory space, many queues can be created. The queues are defined by user-configurable parameters. This flexibility and complexity is necessary in order to create services that offer optimal quality of service and is much better than a restrictive and fixed buffer implementation alternative.

Memory allocation is optimized to guarantee the CBS for each queue. The allocated queue space beyond the CBS that is bounded by the MBS depends on the usage of buffer space and existing guarantees to queues (that is, the CBS). The CBS is defaulted to 8 kB (for 512 byte buffer size) or 18 kB (for 2304 byte buffer size) for all access ingress queues on the 7705 SAR. With a small default queue depth (CBS) allocated for each queue, all services at full scale are guaranteed to have buffers for queuing. The default value would need to be altered to meet the requirements of a specific traffic flow or flows.

3.2.3 Access ingress queuing and scheduling

Traffic management on the 7705 SAR uses a packet-based implementation of the dual leaky bucket model. Each queue has a guaranteed space limited with CBS and a maximum depth limited with MBS. New packets are queued as they arrive. Any packet that causes the MBS to be exceeded is discarded.

The packets in the queue are serviced by two different profiled (rate-based) schedulers, the In-Profile and Out-of-Profile schedulers, where CIR traffic is scheduled before PIR traffic. These two schedulers empty the queue continuously.

For 4-priority scheduling, rate-based schedulers (CIR and PIR) are combined with queue-type schedulers (EXP or BE). For 16-priority scheduling, the rate-based schedulers are combined with the strict priority schedulers (CoS-8 queue first to CoS-1 queue last).



Note: For access ingress and egress, the 16-priority schedulers use additional hardware resources and capabilities, which results in increased throughput.

Access ingress scheduling is supported on the adapter cards and ports listed in the following table. The supported scheduling modes are 4-priority scheduling and 16-priority scheduling. The table shows which scheduling mode each card and port supports at access ingress.

This section also contains information about the following topics:

- [Profiled \(rate-based\) scheduling](#)
- [Queue-type scheduling](#)
- [4-priority scheduling](#)
- [4-priority \(Gen-3\) scheduling](#)
- [16-priority scheduling](#)
- [Ingress queuing and scheduling for BMU traffic](#)

Table 8: Scheduling modes supported by adapter cards and ports at access ingress

Adapter card or port	4-priority	16-priority
8-port Gigabit Ethernet Adapter card	✓	✓
Packet Microwave Adapter card	✓	✓
6-port Ethernet 10Gbps Adapter card ¹	✓	
10-port 1GigE/1-port 10GigE X-Adapter card (in 10-port 1GigE mode)	✓	✓
4-port SAR-H Fast Ethernet module	✓	
6-port SAR-M Ethernet module	✓	✓
Ethernet ports on the 7705 SAR-A	✓	✓
Ethernet ports on the 7705 SAR-Ax	✓	✓
Ethernet ports on the 7705 SAR-H	✓	✓
Ethernet ports on the 7705 SAR-Hc	✓	✓
Ethernet ports on the 7705 SAR-M	✓	✓
Ethernet ports on the 7705 SAR-Wx	✓	✓
Ethernet ports on the 7705 SAR-X ¹	✓	
16-port T1/E1 ASAP Adapter card	✓	
32-port T1/E1 ASAP Adapter card	✓	
2-port OC3/STM1 Channelized Adapter card	✓	

Adapter card or port	4-priority	16-priority
4-port OC3/STM1 Clear Channel Adapter card	✓	
4-port OC3/STM1 / 1-port OC12/STM4 Adapter card	✓	
4-port DS3/E3 Adapter card	✓	
T1/E1 ASAP ports on the 7705 SAR-A	✓	
T1/E1 ASAP ports on the 7705 SAR-M	✓	
TDM ports on the 7705 SAR-X	✓	
12-port Serial Data Interface card	✓	
6-port E&M Adapter card	✓	
6-port FXS Adapter card	✓	
8-port FXO Adapter card	✓	
8-port Voice & Teleprotection card	✓	
8-port C37.94 Teleprotection card	✓	
Integrated Services card	✓	

Note:

1. 4-priority scheduler for Gen-3 adapter card or platform.

3.2.3.1 Profiled (rate-based) scheduling

Each queue is serviced based on the user-configured CIR and PIR values. If the packets that are collected by a scheduler from a queue are flowing at a rate that is less than or equal to the CIR value, the packets are scheduled as in-profile. Packets with a flow rate that exceeds the CIR value but is less than the PIR value are scheduled as out-of-profile. [Figure 16: 4-priority scheduling](#) depicts this behavior by the "In-Prof" and "Out-Prof" labels. This behavior is comparable to the dual leaky bucket implementation in ATM networks. With in-profile and out-of-profile scheduling, traffic that flows at rates up to the traffic contract (that is, CIR) from all the queues is serviced before traffic that flows at rates exceeding the traffic contract. This mode of operation ensures that service-level agreements (SLAs) are honored and traffic that is committed to be transported is switched prior to traffic that exceeds the contract agreement.



Note: A profile is an arithmetical analysis of the rates that are permitted for a particular packet flow; therefore, profiled scheduling may also be called rate-based scheduling.

3.2.3.2 Queue-type scheduling

As well as profiled scheduling described above, queue-type scheduling is supported at access ingress. Queues are divided into two categories, those that are serviced by the Expedited scheduler and those that are serviced by the Best Effort scheduler.

The Expedited scheduler has precedence over the Best Effort scheduler. Therefore, at access ingress, CoS queues that are marked with an Expedited priority are serviced first. Then the Best Effort marked queues are serviced. In a default configuration, the Expedited scheduler services the following CoS queues before the Best Effort scheduler services the rest:

- Expedited scheduler: NC, H1, EF, H2
- Best Effort scheduler: L1, AF, L2, BE

If a packet with an Expedited forwarding class arrives while a Best Effort marked queue is being serviced, the Expedited scheduler takes over and services the Expedited marked CoS queue as soon as the packet from the Best Effort scheduler is serviced.

The schedulers at access ingress in the 7705 SAR service the group of all Expedited queues exhaustively ahead of the group of all Best Effort queues. This means that all expedited queues have to be empty before any packet from a Best Effort queue is serviced.



Note: There is no user configuration for the schedulers. The operation of the schedulers is described for informational purposes. A user can control the mapping of traffic flows based on classification controls, that is, by mapping forwarding classes to as many as eight CoS queues.

The following basic rules apply to the queue-type scheduling of CoS queues:

1. Queues marked for Expedited scheduling are serviced in a round-robin fashion before any queues marked as Best Effort (in a default configuration, these would be queues CoS-8 through CoS-5).
2. These Expedited queues are serviced exhaustively within the round robin. For example, if in a default configuration there are two packets scheduled for service in both CoS-8 and CoS-5, one packet from CoS-8 is serviced, then one packet from CoS-5 is serviced, and then the scheduler returns back to CoS-8, until all the packets are serviced.
3. After the Expedited scheduler has serviced all the packets in the queues marked for Expedited scheduling, the Best Effort scheduler starts serving the queues marked as Best Effort. The same principle described in step 2 is followed, until all the packets in the Best Effort queues are serviced.
4. If a packet arrives at any of the queues marked for Expedited scheduling while the scheduler is servicing a packet from a Best Effort queue, the Best Effort scheduler finishes servicing the current packet and then the Expedited scheduler immediately activates to service the packet in the Expedited queue. If there are no other packets to be serviced in any of the Expedited queues, the Best Effort scheduler resumes servicing the packets in the Best Effort queues. If the queues are configured according to the tables and defaults described in this guide, CoS-4 will be scheduled prior to CoS-1 among queues marked as Best Effort.
5. After one cycle is completed across all the queues marked as Best Effort, the next pass is started until all the packets in all the queues marked as Best Effort are serviced, or a packet arrives to a queue marked as Expedited and is serviced as described in step 2.

3.2.3.3 4-priority scheduling

With 4-priority scheduling, profiled scheduling and queue-type scheduling are combined and the combination is applied to all of the access ingress queues. The profile and queue-type schedulers are combined and applied to the queues to provide maximum flexibility and scalability that meet the stringent QoS requirements of modern network applications. See [Profiled \(rate-based\) scheduling](#) and [Queue-type scheduling](#) for information about these types of scheduling.

Packets with a flow rate that is less than or equal to the CIR value of a queue are scheduled as in-profile. Packets with a flow rate that exceeds the CIR value but is less than the PIR value of a queue are scheduled as out-of-profile.

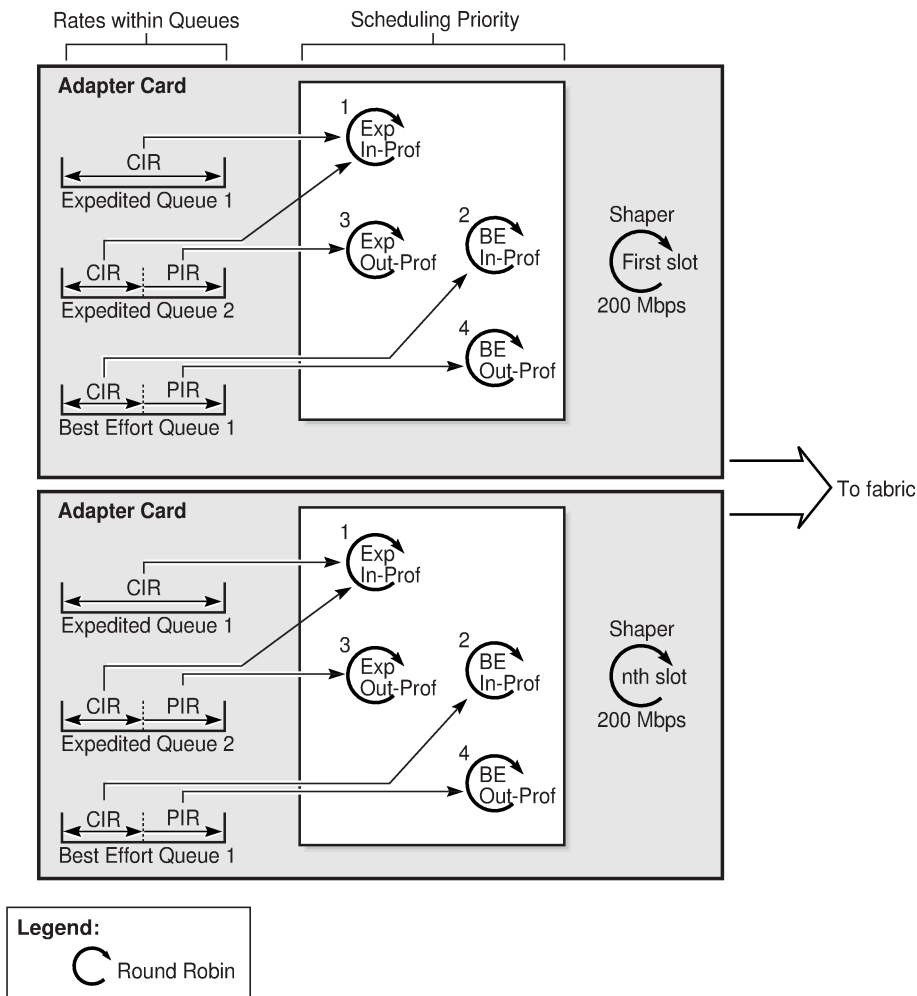
The scheduling cycle for 4-priority scheduling of CoS queues is shown in [Figure 16: 4-priority scheduling](#). The following basic steps apply:

1. In-profile traffic from Expedited queues is serviced in round-robin fashion up to the CIR value. When a queue exceeds its configured CIR value, its state is changed to out-of-profile.
2. When all of the in-profile packets from the Expedited queues are serviced, in-profile packets from Best Effort queues are serviced in a round-robin fashion until the configured CIR value is exceeded. When a queue exceeds its configured CIR value, its state is changed to out-of-profile.
3. When all of the in-profile packets from the Best Effort queues are serviced, out-of-profile packets from Expedited queues are serviced in a round-robin fashion.
4. When all of the out-of-profile packets from the Expedited queues are serviced, the out-of-profile packets from the Best Effort queues are serviced in a round-robin fashion.



Note: If a packet arrives at any of the queues marked for Expedited scheduling while the scheduler is servicing a packet from a Best Effort queue or is servicing an out-of-profile packet, the scheduler finishes servicing the current packet and then returns to the Expedited queues immediately.

Figure 16: 4-priority scheduling



19761

3.2.3.4 4-priority (Gen-3) scheduling

At access ingress, 4-priority scheduling for Gen-3 hardware is the same as 4-priority scheduling for Gen-2 hardware, except that scheduling is done on a per-SAP basis. For more information, see [QoS for Gen-3 adapter cards and platforms](#).

3.2.3.5 16-priority scheduling

For 16-priority scheduling, the rate-based schedulers (CIR and PIR) are combined with the strict priority schedulers (CoS-8 queue first to CoS-1 queue last).

For general information about 16-priority scheduling, see [Network egress 16-priority scheduling](#). Access ingress 16-priority scheduling functions in a similar fashion to network egress 16-priority scheduling.

3.2.3.6 Ingress queuing and scheduling for BMU traffic

The 7705 SAR treats broadcast, multicast, and unknown traffic in the same way as unicast traffic. After being classified, the BMU traffic can be mapped to individual queues in order to be forwarded to the fabric. Classification of unicast and BMU traffic does not differ, which means that BMU traffic that has been classified to a BMU-designated queue can be shaped at its own rate, offering better control and fairer usage of fabric resources. For more information, see [BMU support](#).

3.2.4 Access ingress per-SAP aggregate shapers (access ingress H-QoS)

On the 7705 SAR, H-QoS adds second-tier (or second-level), per-SAP aggregate shapers. As shown in [Figure 17: Access ingress scheduling for 4-priority and 16-priority SAPs \(with per-SAP aggregate shapers\)](#), traffic ingresses at an Ethernet SAP and is classified and mapped to up to eight different CoS queues on a per-ingress SAP basis. The aggregate rate CIR and PIR values are then used to shape the traffic. The conforming loop (aggregate CIR loop) schedules the packets out of the eight CoS queues in strict priority manner (queue priority CIRs followed by queue priority PIRs). If the aggregate CIR is crossed at any time during the scheduling operation, regardless of the per-queue CIR/PIR configuration, then the aggregate conforming loop for the SAP ends and the aggregate non-conforming loop begins.

The aggregate non-conforming loop schedules the packets out of the eight CoS queues in strict priority manner. SAPs sending traffic to the 4-priority scheduler do not have a second-tier per-SAP aggregate shaper unless traffic arbitration is needed, in which case an aggregate CIR for all the 4-priority SAPs can be configured (see [Access ingress per-SAP shapers arbitration](#)). See [Per-SAP aggregate shapers \(H-QoS\) on Gen-2 hardware](#) for general information.

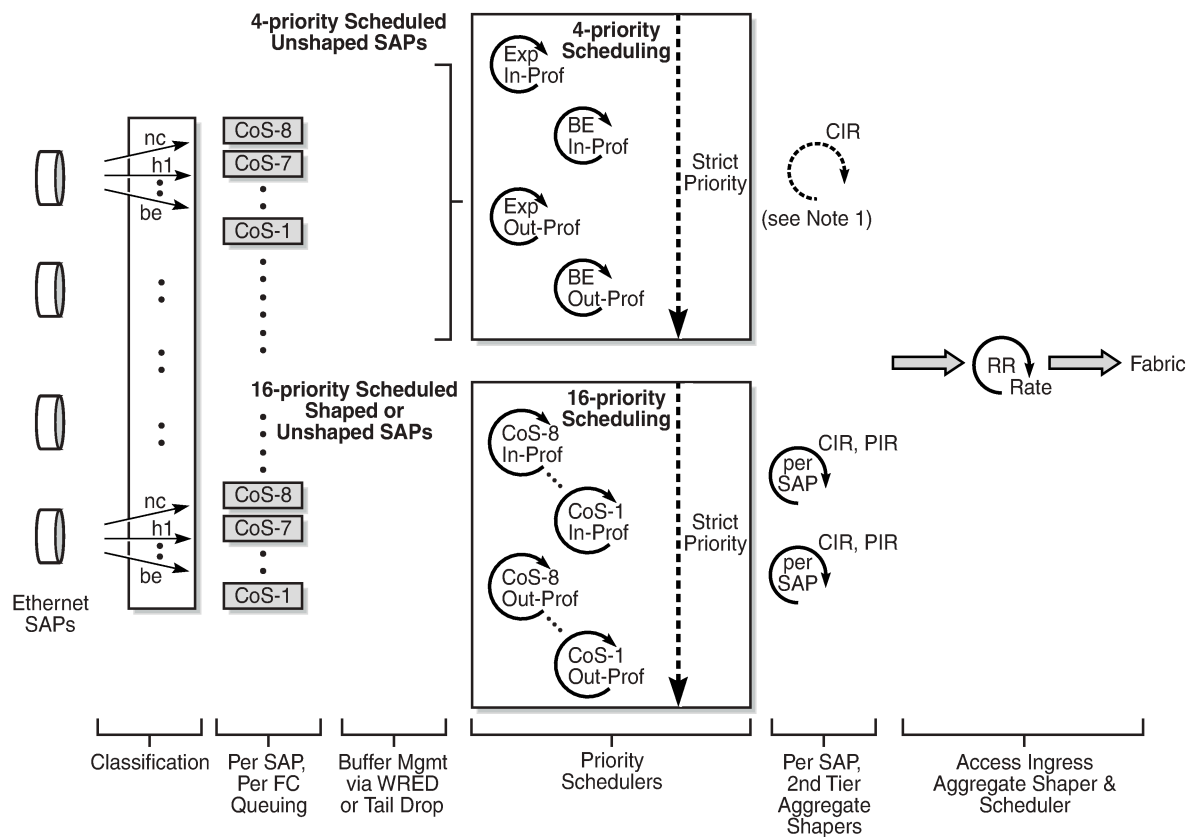
The aggregate rate limit for the per-SAP aggregate shaper is configured in the **service** context, using the **sap>ingress>agg-rate-limit** or **sap>egress>agg-rate-limit** command.

For per-SAP aggregate shaping on Gen-2 adapter cards, the SAP must be scheduled using a 16-priority scheduler.



Note: The default setting for **scheduler-mode** is 4-priority on Gen-2 adapter cards and platforms. The user must toggle the scheduling mode to 16-priority for a SAP before SAP aggregate shaper rates (**agg-rate-limit**) can be configured. Before changing the scheduling mode, the SAP must be shut down.

Figure 17: Access ingress scheduling for 4-priority and 16-priority SAPs (with per-SAP aggregate shapers)



Note 1: Aggregate shaper (CIR) for all the 4-priority unshaped SAPs.

23369

The 16-priority scheduler can be used without setting an aggregate rate limit for the SAP, in which case traffic out of the SAP queues is serviced in strict priority order, the conforming traffic before the non-conforming traffic. Using 16-priority schedulers without a configured per-SAP aggregate shaper (PIR = maximum and CIR = 0 kb/s) may be preferred over 4-priority mode for the following reasons:

- coherent scheduler behavior across SAPs (one scheduler model)
- ease of configuration

As shown in the figure, all the traffic leaving from the shaped SAPs must be serviced using 16-priority scheduling mode.

The SAPs without an aggregate rate limit, which are called unshaped SAPs, can be scheduled using either 4-priority or 16-priority mode as one of the following:

- unshaped SAPs bound to a 4-priority scheduler
- unshaped SAPs bound to a 16-priority scheduler

The arbitration of access ingress traffic leaving the 4-priority and 16-priority schedulers and continuing toward the fabric is described in the following section.

3.2.4.1 Access ingress per-SAP shapers arbitration

The 7705 SAR provides per-SAP aggregate shapers for access ingress SAPs. With this feature, both shaped and unshaped SAPs can coexist on the same adapter card. When switching traffic from shaped and unshaped SAPs to the fabric, arbitration is required.

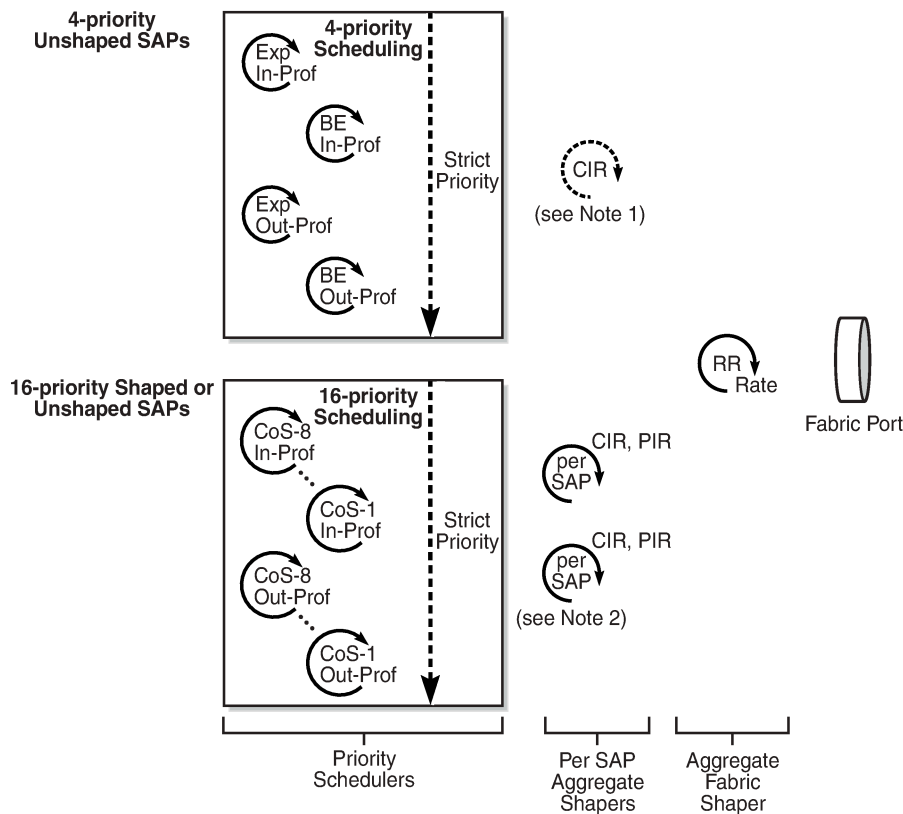
[Figure 18: Access ingress per-SAP arbitration to fabric](#) shows how the 7705 SAR arbitrates traffic to the fabric between 4-priority unshaped SAPs, and 16-priority shaped and unshaped SAPs.

All SAPs support configurable CIR and PIR rates on a per-CoS queue basis (per-queue level). In addition, each 16-priority SAP has its own configurable per-SAP aggregate CIR and PIR rates that operate one level above the per-queue rates.

To allow the 4-priority unshaped SAPs to compete for fabric bandwidth with the aggregate CIR rates of the shaped SAPs, the 4-priority unshaped SAPs (as a group) have their own configurable unshaped SAP aggregate CIR rate, which is configured on the 7705 SAR-8 Shelf V2 and 7705 SAR-18 under the **config>qos>fabric-profile aggregate-mode** context using the **unshaped-sap-cir** parameter. On the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-Wx, the CIR rate is configured in the **config>system>qos>access-ingress-aggregate-rate** context.

The configured CIR and PIR for the 16-priority shaped SAPs dictates committed and uncommitted fabric bandwidth for each of these SAPs. Configuring the **unshaped-sap-cir** parameter for the group (aggregate) of 4-priority unshaped SAPs ensures that the unshaped SAPs will compete for fabric bandwidth with the aggregate CIR rate of the shaped SAPs. Otherwise, the unshaped SAPs would only be able to send traffic into the fabric after the aggregate CIR rates of all the shaped SAPs were serviced. The 16-priority unshaped SAPs are serviced as if they are non-conforming traffic for 16-priority shaped SAPs.

Figure 18: Access ingress per-SAP arbitration to fabric



Note 1: Aggregate shaper for all unshaped 4-priority SAPs.

Note 2: Aggregate shapers for 16-priority SAPs.

23378

The aggregate fabric shaper shown in the figure performs round-robin selection between the 16-priority SAPs (shaped and unshaped) and the 4-priority unshaped SAP aggregate until:

- the aggregate fabric shaper rate is exceeded
- the conforming (CIR) traffic for every 16-priority SAP and the 4-priority unshaped SAP aggregate is exceeded
- the non-conforming traffic for every 16-priority SAP and the 4-priority unshaped SAP aggregate is completed, provided that the aggregate PIR rate is not exceeded



Note: The CLI does not block a fabric profile with an unshaped SAP CIR configuration on the 6-port Ethernet 10Gbps Adapter card or on an ASAP card (16-port T1/E1 ASAP Adapter card or 32-port T1/E1 ASAP Adapter card). However, the unshaped SAP CIR configuration has no effect and is ignored on these cards.

3.2.5 Ingress shaping to fabric (access and network)

After the traffic is scheduled, it must be sent to the fabric interface. In order to avoid congestion in the fabric and ease the effects of possible bursts, a shaper is implemented on each adapter card.

The shapers smooth out any packet bursts and ease the flow of traffic onto the fabric. The shapers use buffer space on the adapter cards and eliminate the need for large ingress buffers in the fabric.

The ingress to-fabric shapers are user-configurable. For the 7705 SAR-8 Shelf V2 and the 7705 SAR-18, the maximum rate depends on a number of factors, including platform, chassis variant, and slot type. See [Configurable ingress shaping to fabric \(access and network\)](#) for details. For the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-Wx, the shapers can operate at a maximum rate of 5 Gb/s. For the 7705 SAR-X, the shapers are not user-configurable. See [Fabric shaping on the fixed platforms \(access and network\)](#) for details.

After the shaping function, all of the traffic is forwarded to the fabric interface in round-robin fashion, one packet at a time, from every access ingress adapter card.



Note: If **per-service-hashing** is not enabled, a 4-byte hash value will be appended to internal overhead for VPLS multicast traffic at ingress. The egress internal hash value is discarded at egress before scheduling. Therefore, shaping rates at access and network ingress and for fabric policies may need to be adjusted accordingly. In addition, the 4-byte internal hash value may be included in any affected statistics counters.

3.2.5.1 BMU support

Fabric shapers support both unicast and multipoint traffic. Multipoint traffic can be any combination of broadcast, multicast, and unknown (BMU) frames. From access ingress to the fabric, BMU traffic is treated as unicast traffic. A single copy of BMU traffic is handed off to the fabric, where it is replicated and sent to all potential destination adapter cards.

3.2.5.1.1 Aggregate mode BMU support

An aggregate mode shaper provides a single aggregate shaping rate. The rate defines the maximum bandwidth that an adapter card can switch through its fabric interface at any given time. The rate is a bulk value and is independent of the destination or the type of traffic. For example, in aggregate mode, an ingress adapter card may use the full rate to communicate with a single destination adapter card, or it may use the same rate to communicate with multiple egress adapter cards.

Aggregate mode and the aggregate rate apply to fabric shapers that handle combined unicast/BMU traffic, unicast-only traffic, or BMU-only traffic. One aggregate rate sets the rate on all adapter cards. The proportional distribution between unicast and BMU traffic can be fine-tuned using queue-level schedulers, while the to-fabric shaper imposes a maximum rate that ensures fairness on the fabric for traffic from all adapter cards.

When services (IES, VPRN, and VPLS) are enabled, the fabric profile mode for access ingress should be set to aggregate mode.

3.2.5.1.2 Destination mode BMU support

Destination mode offers granular to-fabric shaping rates on a per-destination adapter card basis. While destination mode offers more flexibility and gives more control than aggregate mode, it also requires a greater understanding of network topology and flow characteristics under conditions such as node failures and link, adapter card, or port failures.

In a destination mode fabric profile, the unicast traffic and BMU traffic are always shaped separately.

For unicast traffic, individual destination rates can be configured on each adapter card. For BMU traffic, one multipoint rate sets the rate on all adapter cards. Fairness among different BMU flows is ensured by tuning the QoS queues associated with the port.

3.2.5.2 LAG SAP support (access only)

Fabric shapers support access ingress traffic being switched from a SAP to another SAP residing on a port that is part of a link aggregation group (LAG). Either the aggregate mode or destination mode can be used for fabric shaping.

When the aggregate mode is used, one aggregate rate sets the rate on all adapter cards. When the destination mode is used, the multipoint shaper is used to set the fabric shaping rate for traffic switched to a LAG SAP.



Note: Even though the multipoint shaper is used to set the fabric shaping rate for traffic switched to a LAG SAP, it is the per-destination unicast counters that are incremented to show the fabric statistics rather than the multipoint counter. Only the fabric statistics of the active port of the LAG are incremented, not the standby port.

3.2.6 Configurable ingress shaping to fabric (access and network)

The use of fabric profiles allows the ingress (to the fabric) shapers to be user-configurable for access ingress and network ingress traffic.

For the 7705 SAR-8 Shelf V2 and 7705 SAR-18, the maximum rates are:

- 2.5 Gb/s for the 7705 SAR-8 Shelf V2 (all 6 MDA slots)
- 10 Gb/s for the 7705 SAR-8 Shelf V2 (MDA slots 1 and 2)
- 1 Gb/s or 2.5 Gb/s for the 7705 SAR-18 (12 MDA slots)
- 10 Gb/s for the 7705 SAR-18 (4 XMDA slots)

For information about fabric shapers on the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-X, see [Fabric shaping on the fixed platforms \(access and network\)](#).

By allowing a rate of 1 Gb/s or higher to be configured from any adapter card to the fabric, the fabric may become congested. Therefore, the collection and display of fabric statistics are provided. These statistics report about the fabric traffic flow and potential discards. See the 7705 SAR Interface Configuration Guide, "Configuring Adapter Card Fabric Statistics", "Configuration Command Reference", and "Show, Monitor, Clear, and Debug Command Reference" for information about how to configure, show, and monitor fabric statistics on an adapter card.

The ingress buffers for a card are much larger than the ingress buffers for the fabric; therefore, it is advantageous to use the larger card buffers for ingress shaping. In order to use the ingress card buffers

and have much more granular control over traffic, two fabric profile modes are supported, per-destination mode and aggregate mode. Both modes offer shaping toward the fabric from an adapter card, but per-destination shapers offer the maximum flexibility by precisely controlling the amount of traffic to each destination card at a user-defined rate. Aggregate mode is used for simpler deployments, where the amount of traffic flowing to a destination adapter card is not controlled.

The default mode of operation for the 7705 SAR is set to aggregate, and the fixed aggregate rate of 200 Mb/s is set for both access ingress and network ingress traffic. Therefore, in a default configuration, each adapter card can switch up to 200 Mb/s of access ingress and network ingress traffic toward the fabric.

All the switched traffic can be destined for a single adapter card or it can be spread among multiple adapter cards. For higher-bandwidth applications, a network traffic analysis is recommended to determine which shaper rates would best suit the application and traffic patterns of a particular environment.

The to-fabric shapers are provided on the 7705 SAR to ensure adequate use of ingress buffers in case of congestion. With the ingress shapers, the large ingress card buffers can be configured to absorb bursty traffic and pace the traffic for better use of resources.

For example, if the average access ingress traffic bandwidth for an adapter card is 400 Mb/s and the peak bandwidth is 800 Mb/s, the rate of the to-fabric shapers can be configured to be 400 Mb/s. This allows the bursty ingress traffic to be paced by absorbing the bursty traffic after being shaped at 400 Mb/s. The initial burst is absorbed at the adapter card where the bursty traffic ingresses the 7705 SAR. The ingress buffers are used to absorb the burst and the fabric buffers are not exhausted by any single adapter card. The same example applies to network ingress traffic.

The following table summarizes the different capabilities offered by the two modes.

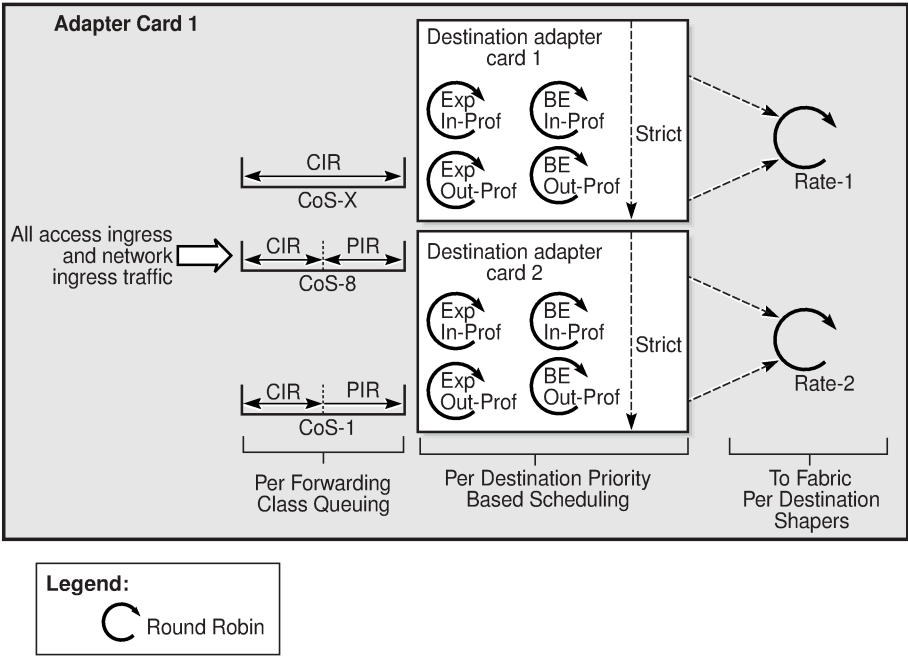
Table 9: Fabric profile mode options and capabilities

Capability	Per-destination mode	Aggregate mode
Access ingress to-fabric shapers	✓	✓
Network ingress to-fabric shapers	✓	✓
Individual shaping from an ingress card toward each destination card based on a user-defined rate	✓	
Aggregate/bulk sum shaping regardless of destination from an ingress card		✓

Figure 19: Fabric shapers in per-destination mode and Figure 20: Fabric shapers in aggregate mode illustrate the functionality of fabric shapers in per-destination mode and aggregate mode, respectively.

In the following figure, after the per-destination prioritization and scheduling takes place as described in previous sections in this chapter, the per-destination adapter card shapers take effect. With per-destination shapers, the maximum amount of bandwidth that each destination adapter card can receive from the fabric can be controlled. For example, the maximum amount of bandwidth that adapter card 1 can switch to the remaining adapter cards, as well as the amount of bandwidth switched back to adapter card 1, can be configured at a set rate.

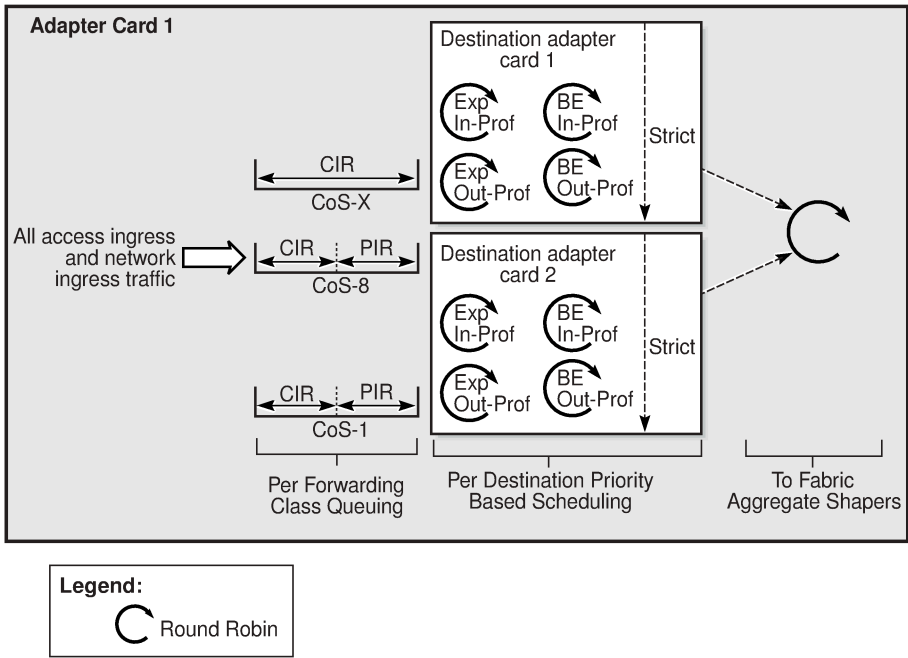
Figure 19: Fabric shapers in per-destination mode



20174

The following figure illustrates the functionality of fabric shapers in aggregate mode. After the policing, classification, queuing and per-destination based priority queuing takes place as described in previous sections in this chapter, the aggregate mode adapter card shapers take effect. In aggregate mode, the aggregate of all the access ingress and network ingress traffic is shaped at a user-configured rate regardless of the destination adapter card.

Figure 20: Fabric shapers in aggregate mode



20175

Mixing different fabric shaper modes within the same chassis and on the same adapter card is not recommended; however, it is supported. As an example, an 8-port Gigabit Ethernet Adapter card in a 7705 SAR-18 can be configured for aggregate mode for access ingress and for per-destination mode for network ingress. The same chassis can also contain an adapter card (for example, the 32-port T1/E1 ASAP Adapter card) that is configured for per-destination mode for all traffic. This setup is shown in the following example.

	MDA	Access fabric mode	Network profile mode
1/1	a8-1gb-v3-sfp	Destination	Destination
1/2	a8-1gb-sfp	Aggregate	Destination
1/3	a4-oc3	Destination	Destination
1/4	a32-chds1v2	Destination	Destination
1/X1	x-10GigE-v2	Aggregate	Destination



- Note:**
- Gen-2 and Gen-3 adapter cards only support aggregate mode fabric shapers for access ingress traffic, regardless of the service types configured.
 - If multipoint services such as IES, VPRN, and VPLS are running on an adapter card, only aggregate mode fabric profile can be configured for the card at access ingress.

3.2.7 Fabric shaping on the fixed platforms (access and network)

The 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, and 7705 SAR-Wx support user-configurable fabric shapers at rates of up to 5 Gb/s for access ingress and network ingress traffic. The fabric interface on these nodes is a shared resource between both traffic types, and one buffer pool serves all MDAs (ports).

These nodes do not support fabric profiles; instead, they have a single aggregate rate limit for restricting access traffic into the fabric and a single aggregate rate limit for restricting network traffic into the fabric. These limits apply to all MDAs. Both access ingress and network ingress traffic can be configured to shaping rates of between 1 kb/s and 5 Gb/s. The default rate for access ingress traffic is 500 Mb/s. The default rate for network ingress traffic is 2 Gb/s. Statistics can be viewed for aggregate access and network traffic flow through the fabric and possible discards.

The 7705 SAR-X fabric shaper rate is not configurable for access ingress or network ingress traffic, and is set to the maximum rate for the platform. There are three buffer pools on the 7705 SAR-X, one for each MDA (block of ports).

3.3 Traffic flow across the fabric

The 7705 SAR uses an Ethernet-based fabric. Each packet that is sent to the fabric is equipped with a fabric header that contains its specific CoS requirement. Because all of the packets switched across the fabric are already classified, queued, scheduled and marked according to the required QoS parameters, each of these packets has been passed through the Traffic Management (TM) block on an adapter card, or the Control and Switching Module (CSM) in the case of control packets. Therefore, each packet arrives at the fabric having been already scheduled for optimal flow. The function of the fabric is to switch each packet through to the appropriate destination adapter card, or CSM in the case of control packets, in an efficient manner.

Because the traffic is shaped at a certain rate by the ingress adapter card (that is, bursts are smoothed by the traffic management function), minimal buffering should be needed on the switch fabric. However, the buffer space allocation and usage is in accordance with the priorities at the ingress adapter card. As is the case with schedulers at the adapter cards, there are two priorities supported on the switch fabric. The switch fabric serves the traffic in the following priority order:

1. Expedited
2. Best Effort



Note:

- The switch fabric does not support profile scheduling.
- Because the fabric has a limited buffer space, it is possible for tail drop to occur. Tail drop discards any packet that exceeds the maximum buffer space allocation. The shaping that is performed on the adapter cards helps to prevent or minimize congestion.

3.4 Network egress

This section contains the following topics for traffic flow in the network egress direction:

- [BMU traffic at network egress](#)
- [Network egress queuing aggregation](#)
- [Network egress scheduling](#)
- [Network egress shaping](#)
- [Network egress shaping for hybrid ports](#)
- [Network egress per-VLAN shapers](#)
- [Network egress marking and re-marking](#)

3.4.1 BMU traffic at network egress

BMU traffic at network egress is handled in the same way as unicast traffic in terms of scheduling, queuing, or port-level shaping. Both unicast and BMU traffic are mapped to queues as per the FC markings. Traffic from these queues, whether unicast or BMU, is scheduled according to user-configured rates. Port-level shapers treat all the queues identically, regardless of traffic type.

3.4.2 Network egress queuing aggregation

After traffic is switched through the fabric from one or several access ingress adapter cards to a network egress adapter card, queuing-level aggregation on a per-forwarding-class basis is performed on all of the received packets.

An adapter card that is used for network egress can receive—and will likely receive—packets from multiple adapter cards that are configured for access ingress operations, and from the CSM. Adapter cards that are configured for network access allow user configuration of queues and the association of forwarding classes to the queues. These are the same configuration principles that are used for adapter cards that are configured for access ingress connectivity. Like access ingress, more than one forwarding class can share the same queue.

Aggregation of different forwarding classes under queues takes place for each bundle or port. If a port is a member of a bundle, such as a Multilink Point-to-Point Protocol (MLPPP) bundle, then the aggregation and queuing is implemented for the entire bundle. If a port is a standalone port, that is, not a member of bundle, then the queuing takes place for the port.

3.4.2.1 Network egress per-VLAN queuing

Network Ethernet ports support network egress per-VLAN (per-interface) shapers with eight CoS queues per VLAN, which is an extension to the eight CoS queues per port shared by all unshaped VLANs. Eight unique per-VLAN CoS queues are created for each VLAN when the VLAN shaper is enabled. These per-VLAN CoS queues are separate from the eight unshaped VLAN queues. The eight CoS queues that are shared by all the remaining unshaped VLANs are referred to as unshaped VLAN CoS queues. VLAN shapers are enabled when the **queue-policy** command is used to assign a network queue policy to the interface.

For details on per-VLAN network egress queuing and scheduling, see [Per-VLAN network egress shapers](#).

3.4.3 Network egress scheduling

Network egress scheduling is supported on the adapter cards and ports listed in the following table. The supported scheduling modes are 4-priority and 16-priority. The table shows which scheduling mode each card and port supports at network egress.

This section also contains information about the following topics:

- [Network egress 4-priority scheduling](#)
- [Network egress 4-priority \(Gen-3\) scheduling](#)
- [Network egress 16-priority scheduling](#)

Table 10: Scheduling modes supported by adapter cards and ports at network egress

Adapter card or port	4-priority	16-priority
8-port Gigabit Ethernet Adapter card		✓
Packet Microwave Adapter card		✓
2-port 10GigE (Ethernet) Adapter card/module		✓
6-port Ethernet 10Gbps Adapter card ¹	✓	
10-port 1GigE/1-port 10GigE X-Adapter card		✓
4-port SAR-H Fast Ethernet module		✓
6-port SAR-M Ethernet module		✓
Ethernet ports on the 7705 SAR-A		✓
Ethernet ports on the 7705 SAR-Ax		✓
Ethernet ports on the 7705 SAR-H		✓
Ethernet ports on the 7705 SAR-Hc		✓
Ethernet ports on the 7705 SAR-M		✓
Ethernet ports on the 7705 SAR-Wx		✓
Ethernet ports on the 7705 SAR-X ¹	✓	
16-port T1/E1 ASAP Adapter card	✓	
32-port T1/E1 ASAP Adapter card	✓	
2-port OC3/STM1 Channelized Adapter card	✓	
4-port OC3/STM1 / 1-port OC12/STM4 Adapter card	✓	
4-port OC3/STM1 Clear Channel Adapter card	✓	

Adapter card or port	4-priority	16-priority
4-port DS3/E3 Adapter card	✓	
T1/E1 ASAP ports on the 7705 SAR-A	✓	
T1/E1 ASAP ports on the 7705 SAR-M	✓	
TDM ports on the 7705 SAR-X	✓	

Note:

1. 4-priority scheduler for Gen-3 adapter card or platform.

3.4.3.1 Network egress 4-priority scheduling

The implementation of network egress scheduling on the cards and ports listed in [Table 10: Scheduling modes supported by adapter cards and ports at network egress](#) under "4-Priority" is very similar to the scheduling mechanisms used for adapter cards that are configured for access ingress traffic. 4-priority scheduling is a combination of queue-type scheduling (Expedited vs. Best-effort scheduling) and profiled scheduling (rate-based scheduling).



Note: The encapsulation type must be ppp-auto for PPP/MLPPP bundles on the following:

- T1/E1 ports on the 7705 SAR-A
- T1/E1 ports on the 7705 SAR-M
- T1/E1 ports on the 7705 SAR-X
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- 2-port OC3/STM1 Channelized Adapter card
- 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card
- T1/E1 ports on the 4-port T1/E1 and RS-232 Combination module (on 7705 SAR-H)

Packets less than or up to the CIR are scheduled as in-profile. Packets that arrive at rates greater than the CIR, but less than the PIR, are scheduled as out-of-profile. In-profile traffic is exhaustively transmitted from the queues before out-of-profile traffic is transmitted. That is, all of the in-profile packets must be transmitted before any out-of-profile packets are transmitted. In addition, Expedited queues are always scheduled before Best Effort queues.

The default configuration of scheduling CoS queues provides a logical and consistent means to manage the traffic priorities. The default configuration is as follows:

- CoS-8 to CoS-5 Expedited in-profile
- CoS-4 to CoS-1 Best Effort in-profile
- CoS-8 to CoS-5 Expedited out-of-profile
- CoS-4 to CoS-1 Best Effort out-of-profile



Note: Default configuration means that the queues are configured according to the tables and defaults described in this guide. Customers can configure the queues differently.

The order shown below is maintained when scheduling the traffic on the adapter card's network ports. A strict priority is applied between the four schedulers, and all four schedulers are exhaustive:

- Expedited in-profile traffic
- Best Effort in-profile traffic
- Expedited out-of-profile traffic
- Best Effort out-of-profile traffic

3.4.3.2 Network egress 4-priority (Gen-3) scheduling

The adapter cards and ports that support 4-priority scheduling for network egress traffic on Gen-3 hardware are identified in [Table 10: Scheduling modes supported by adapter cards and ports at network egress](#). This type of scheduling takes into consideration the traffic's profile type and the CoS queue priority. It also uses priority information to apply backpressure to lower-level CoS queues. See [QoS for Gen-3 adapter cards and platforms](#) for details.

3.4.3.3 Network egress 16-priority scheduling

The adapter cards and ports that support 16-priority scheduling for network egress traffic are listed in [Table 10: Scheduling modes supported by adapter cards and ports at network egress](#). This type of scheduling takes into consideration the traffic's profile type and the priority of the CoS queue that the traffic is coming from.

Packets less than or up to the CIR are scheduled as in-profile. Packets that arrive at rates greater than the CIR, but less than the PIR, are scheduled as out-of-profile. Eight CoS queues in total are available for packets to go through.

In-profile traffic is exhaustively transmitted from the queues, starting with the highest-priority CoS queue. A strict priority is applied between the eight CoS queues. If a packet arrives at a queue of higher priority than the one being serviced, the scheduler services the packet at the higher-priority queue as soon as it finishes servicing the current packet.

When all the in-profile traffic is transmitted, the out-of-profile traffic is transmitted, still maintaining priority of the queues. If an in-profile packet arrives and the scheduler is servicing an out-of-profile packet, the scheduler finishes servicing the out-of-profile packet and then immediately services the in-profile packet.

The order of priority in the default configuration is as follows:

- CoS-8 in-profile traffic
- CoS-7 in-profile traffic
- CoS-6 in-profile traffic
- CoS-5 in-profile traffic
- CoS-4 in-profile traffic
- CoS-3 in-profile traffic
- CoS-2 in-profile traffic
- CoS-1 in-profile traffic

- CoS-8 out-of-profile traffic
- CoS-7 out-of-profile traffic
- CoS-6 out-of-profile traffic
- CoS-5 out-of-profile traffic
- CoS-4 out-of-profile traffic
- CoS-3 out-of-profile traffic
- CoS-2 out-of-profile traffic
- CoS-1 out-of-profile traffic



Note: Default configuration means that the queues are configured according to the tables and defaults described in this guide. Customers can configure the queues differently.

3.4.4 Network egress shaping

All the network egress traffic is shaped at the bundle or interface rate. An interface may not necessarily correspond directly to a port, and an interface could be a sub-channel of a port. As an example, Fast Ethernet could be the choice of network egress, but the leased bandwidth could still be a fraction of the port speed. In this case, it is possible to shape at the interface rate of 15 Mb/s, for example.

The same also applies to MLPPP bundles. The shaping takes place per MLPPP bundle, and the traffic is shaped at the aggregate rate of the MLPPP bundle.

3.4.5 Network egress shaping for hybrid ports

Hybrid ports use a third-tier, dual-rate aggregate shaper to provide arbitration between the bulk of access and network egress traffic flows. For details, see [QoS for hybrid ports on Gen-2 hardware](#).

3.4.6 Network egress per-VLAN shapers

Network egress VLAN traffic uses second-tier (or second-level), per-VLAN shapers to prepare network egress traffic for arbitration with the aggregate of the unshaped VLAN shaper. All the shaped VLAN shapers are arbitrated with one unshaped VLAN shaper.



Note: This section applies to Gen-2 adapter cards and platforms. For information about per-VLAN shapers for Gen-3 adapter cards and platforms, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X, see [Figure 13: 4-priority scheduling at network egress \(Gen-3 hardware\) on a network port](#) in the [QoS for Gen-3 adapter cards and platforms](#) section.


As shown in the following figure, traffic from the fabric flows to one or more VLANs where it is classified and mapped to up to eight different CoS queues on a per-VLAN basis. The VLANs can be shaped or unshaped. Each shaped VLAN has its own set of CoS queues. The aggregate of unshaped VLANs uses the same set of CoS queues (that is, one set of queues for all unshaped VLANs).

For more information, see [Per-VLAN network egress shapers](#) and [Shaped and unshaped VLANs](#).



Note: Because of space limitations in the figure, the second-tier, per-VLAN aggregate shapers are represented as a single loop containing the label “per VLAN”, even though they are dual-rate shapers similar to the third-tier network aggregate shaper.

Figure 21: Network egress shaped and unshaped VLAN queuing and scheduling



Note 1: Each per-VLAN aggregate shaper is a dual rate shaper similar to the 3rd tier dual rate shaper

24354

Because the per-VLAN shapers are dual-rate shapers, their aggregate rate CIR and PIR values shape the traffic, as follows:

- The conforming, in-profile loop (aggregate CIR loop) schedules the packets out of the eight CoS queues in strict priority manner (queue priority CIRs followed by queue priority PIRs).
- If the aggregate CIR is crossed at any time during the scheduling operation, regardless of the per-queue CIR/PIR configuration, then the aggregate conforming loop for the VLAN ends and the aggregate non-conforming loop (out-of-profile) begins.
- The aggregate non-conforming loop schedules the packets out of the eight CoS queues in strict priority manner.

A shaped VLAN configured with default aggregate rate limits (PIR = maximum and CIR = 0 kb/s) is equivalent to an unshaped VLAN except that its traffic flows through a per-VLAN shaper instead of getting combined with the bulk (aggregate) of the unshaped VLANs. Using a shaped VLAN in this way (default rate limits) may be preferred over using an unshaped VLAN for the following reasons:

- coherent scheduler behavior across VLANs (that is, the use of only one scheduler model)
- ease of configuration
- higher throughput, as each shaped VLAN gets to transmit one packet at each pass of the out-of-profile scheduler as opposed to one packet from the aggregate of unshaped VLAN queues

The arbitration of shaped and unshaped VLAN traffic at the third-tier shaper is described in the following section.

3.4.6.1 Network egress per-VLAN shapers arbitration

For shaped VLANs, the configured CIR and PIR limits dictate committed and uncommitted port bandwidth for each of these VLANs. To ensure that the bulk (aggregate) of unshaped VLANs can compete for port bandwidth with the aggregate of CIR rates for the shaped VLANs, the unshaped VLANs (as a group) have their own aggregate CIR rate, which is configured using the **unshaped-if-cir** command (under the **config>port> ethernet>network>egress** context). Otherwise, without their own aggregate CIR rate, the unshaped VLANs are only able to send traffic into the port after the aggregate CIR rates of all the shaped VLANs are serviced. Shaped VLANs using default aggregate rate limits (PIR = maximum and CIR = 0 kb/s) are serviced as if they are non-conforming traffic for shaped VLANs.



Note: This section applies to Gen-2 adapter cards. For information about per-VLAN shapers arbitration for Gen-3 adapter cards and platforms, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X, see [Figure 13: 4-priority scheduling at network egress \(Gen-3 hardware\) on a network port](#) in the [QoS for Gen-3 adapter cards and platforms](#) section.

Referring to [Figure 21: Network egress shaped and unshaped VLAN queuing and scheduling](#), at the port shaper, conforming (CIR) traffic has priority over non-conforming traffic. The arbitration between the shaped VLANs and unshaped VLANs is handled in the following priority order:

- committed traffic: the per-VLAN committed rate (CIR) for shaped VLANs as set by the **agg-rate-limit** command, and the aggregate committed rate for all the unshaped VLANs as set by the **unshaped-if-cir** command
- uncommitted traffic: the per-VLAN uncommitted rate (PIR) for shaped VLANs as set by the **agg-rate-limit** command, and aggregate uncommitted rate for all the unshaped VLANs as set by the **unshaped-if-cir** command

3.4.7 Network egress marking and re-marking

The EXP bit settings can be marked at network egress. The EXP bit markings of the forwarding class are used for this purpose. The tunnel and pseudowire EXP bits are marked to the forwarding class value.

The default network egress QoS marking settings are listed in [Table 24: Default network QoS policy egress marking](#).

3.4.7.1 Network egress marking and re-marking on Ethernet ports

For MPLS tunnels, if network egress Ethernet ports are used, dot1p bit marking can be enabled in conjunction with EXP bit marking. In this case, the tunnel and pseudowire EXP bits do not have to be the same as the dot1p bits.

For GRE and IP tunnels, dot1p marking and pseudowire EXP marking can be enabled, and DSCP marking can also be enabled.

Network egress dot1p is supported for Ethernet frames, which can carry IPv4, IPv6, or MPLS packets. EXP re-marking is supported for MPLS packets.



Note: IP traffic from network interfaces is always trusted. There is no re-marking performed on network egress for global routing table (GRT) forwarded IP traffic (with the exception of GRE and IPSec tunnels). This also applies to IES interface traffic that is forwarded to the network egress interface.

3.5 Network ingress

This section contains the following topics for traffic flow in the network ingress direction:

- [Network ingress classification](#)
- [Network ingress queuing](#)
- [Network ingress scheduling](#)
- [Network ingress shaping to fabric](#)
- [Configurable network ingress shaping to fabric](#)
- [Network fabric shaping on the fixed platforms](#)

3.5.1 Network ingress classification

Network ingress traffic originates from a network egress port located on another interworking device, such as a 7750 Service Router or another 7705 SAR, and flows from the network toward the fabric in the 7705 SAR.

The ingress MPLS packets can be mapped to forwarding classes based on EXP bits that are part of the headers in the MPLS packets. These EXP bits are used across the network to ensure an end-to-end network-wide QoS offering. With pseudowire services, there are two labels, one for the MPLS tunnel and one for the pseudowire. Mapping is performed using the EXP values from the outer tunnel MPLS label. This ensures that the EXP bit settings, which may have been altered along the path by the tandem label switch routers (LSRs), are used to identify the forwarding class of the encapsulated traffic.

Ingress GRE and IP packets are mapped to forwarding classes based on DSCP bit settings of the IP header. GRE tunnels are not supported for IPv6; therefore, DSCP bit classification of GRE packets is only supported for IPv4. DSCP bit classification of IP packets is supported for both IPv4 and IPv6.

Untrusted traffic uses multi-field classification (MFC), where the traffic is classified based on any IP criteria currently supported by the 7705 SAR filter policies; for example, source and destination IP address, source and destination port, whether the packet is fragmented, ICMP code, and TCP state. For information about MFC, see the 7705 SAR Router Configuration Guide, "Multi-field classification" and "IP, MAC, and VLAN filter entry commands".

3.5.1.1 Network ingress tunnel QoS override

To simplify QoS management through the network core, some operators aggregate multiple forwarding classes of traffic at the ingress LER or PE and use two or three QoS markings instead of the eight different

QoS markings that a customer device may be using to dictate QoS treatment. However, to ensure the end-to-end QoS enforcement required by the customer, the aggregated markings must be mapped back to their original forwarding classes at the egress LER (eLER) or PE.

For IP traffic (including IPSec packets) riding over MPLS or GRE tunnels that will be routed to the base router, a VPRN interface, or an IES interface at the tunnel termination point (the eLER), the 7705 SAR can be configured to ignore the EXP/DSCP bits in the tunnel header when the packets arrive at the eLER. Instead, classification is based on the inner IP header, which is essentially the customer IP packet header. This configuration is done using the **ler-use-dscp** command.

When the command is enabled on an ingress network IP interface, the IP interface will ignore the tunnel's QoS mapping and will derive the internal forwarding class and associated profile state based on the DSCP values of the IP header ToS field rather than on the network QoS policy defined on the IP interface. This function is useful when the mapping for the tunnel QoS marking does not completely reflect the required QoS handling for the IP packet. The command applies only on the eLER where the tunnel or service is terminated and the next header in the packet is IP.

3.5.2 Network ingress queuing

Network ingress traffic can be classified in up to eight different forwarding classes, which are served by 16 queues (eight queues for unicast traffic and eight queues for multicast (BMU) traffic). Each queue serves at least one of the eight forwarding classes that are identified by the incoming EXP bits. These queues are automatically created by the 7705 SAR. The following table shows the default network QoS policy for the 16 CoS queues.

The value for CBS and MBS is a percentage of the size of the buffer pool for the adapter card. MBS can be shared across queues, which allows overbooking to occur.

Table 11: Default network ingress QoS policy

Queue /FC	CIR (%)	PIR (%)	CBS (%)	MBS (%)
Queue-1/BE	0	100	0.1	5
Queue-2/L2	25	100	0.25	5
Queue-3/AF	25	100	0.75	5
Queue-4/L1	25	100	0.25	2.5
Queue-5/H2	100	100	0.75	5
Queue-6/EF	100	100	0.75	5
Queue-7/H1	10	100	0.25	2.5
Queue-8/NC	10	100	0.25	2.5
Queue-9/BE	0	100	0.1	5
Queue-10/L2	5	100	0.1	5
Queue-11/AF	5	100	0.1	5

Queue /FC	CIR (%)	PIR (%)	CBS (%)	MBS (%)
Queue-12/L1	5	100	0.1	2.5
Queue-13/H2	100	100	0.1	5
Queue-14/EF	100	100	0.1	5
Queue-15/H1	10	100	0.1	2.5
Queue-16/NC	10	100	0.1	2.5

3.5.2.1 Network ingress queuing for BMU traffic

At network ingress, broadcast, multicast, and unknown (BMU) traffic identified using DSCP and/or EXP (also known as LSP TC) is mapped to a forwarding class (FC). Because BMU traffic is considered to be multipoint traffic, the queue hosting BMU traffic must be configured with the **multipoint** keyword. Queues 9 through 16 support multipoint traffic (see [Table 11: Default network ingress QoS policy](#)). For any adapter card hosting any number of network ports, up to 16 queues can be configured to host 8 unicast and 8 multicast queues.

Similar to unicast queues, BMU queues require configuration of:

- queue depth (committed and maximum)
- scheduled rate (committed and peak)

In addition, as is the case for unicast queues, all other queue-based congestion management techniques apply to multipoint queues.

The benefits of using multipoint queues occur when the to-fabric shapers begin scheduling traffic toward the destination line card. To-fabric shapers can be configured for **aggregate** or **per-destination** mode. For more information, see [BMU support](#).

3.5.3 Network ingress scheduling

Network ingress scheduling is supported on the adapter cards and ports listed in the following table. The supported scheduling modes are 4-priority and 16-priority. The table shows which scheduling mode each card and port supports at network ingress.

This section also contains information about the following topics:

- [Network ingress 4-priority scheduling](#)
- [Network ingress 4-priority \(Gen-3\) scheduling](#)
- [Network ingress 16-priority scheduling](#)

Table 12: Scheduling modes supported by adapter cards and ports at network ingress

Adapter card or port	4-priority	16-priority
8-port Gigabit Ethernet Adapter card		✓
Packet Microwave Adapter card		✓

Adapter card or port	4-priority	16-priority
2-port 10GigE (Ethernet) Adapter card/module		✓
6-port Ethernet 10Gbps Adapter card ¹	✓	
10-port 1GigE/1-port 10GigE X-Adapter card		✓
4-port SAR-H Fast Ethernet module		✓
6-port SAR-M Ethernet module		✓
Ethernet ports on the 7705 SAR-A		✓
Ethernet ports on the 7705 SAR-Ax		✓
Ethernet ports on the 7705 SAR-M		✓
Ethernet ports on the 7705 SAR-H		✓
Ethernet ports on the 7705 SAR-Hc		✓
Ethernet ports on the 7705 SAR-Wx		✓
Ethernet ports on the 7705 SAR-X ¹	✓	
16-port T1/E1 ASAP Adapter card	✓	
32-port T1/E1 ASAP Adapter card	✓	
2-port OC3/STM1 Channelized Adapter card	✓	
4-port OC3/STM1 Clear Channel Adapter card	✓	
4-port OC3/STM1 / 1-port OC12/STM4 Adapter card	✓	
4-port DS3/E3 Adapter card	✓	
T1/E1 ASAP ports on the 7705 SAR-A	✓	
T1/E1 ASAP ports on the 7705 SAR-M	✓	
TDM ports on the 7705 SAR-X	✓	

Note:

1. 4-priority scheduler for Gen-3 adapter card or platform.

3.5.3.1 Network ingress 4-priority scheduling

The adapter cards listed in [Table 12: Scheduling modes supported by adapter cards and ports at network ingress](#) under "4-Priority" can receive network ingress traffic. One or more ports on the card are configured for PPP/MLPPP for this purpose.

The implementation of network ingress scheduling on the cards listed in the table under "4-Priority" is very similar to the scheduling mechanisms used for adapter cards that are configured for access ingress traffic. That is, 4-priority scheduling is used (queue-type scheduling combined with profiled scheduling).



Note: The encapsulation type must be ppp-auto for PPP/MLPPP bundles on the following:

- T1/E1 ports on the 7705 SAR-A
- T1/E1 ports on the 7705 SAR-M
- T1/E1 ports on the 7705 SAR-X
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- 2-port OC3/STM1 Channelized Adapter card
- 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card
- T1/E1 ports on the 4-port T1/E1 and RS-232 Combination module (on 7705 SAR-H)

The adapter cards provide sets of eight queues for incoming traffic: 7 sets of queues for the 7705 SAR-8 Shelf V2 and 17 sets of queues for the 7705 SAR-18. Each set of queues is specific to a destination adapter card. For the 7705 SAR-8 Shelf V2 and 7705 SAR-18 (respectively), 6 and 16 sets of queues are automatically created for each access egress adapter card, plus 1 set of queues for multicast traffic.

There is one additional set of queues for slow-path (control) traffic destined for the CSMs.

The individual queues within each set of queues provide buffer space for traffic isolation based on the CoS values being applied (from the received EXP bits).

All of the network ingress ports of the adapter card share the same sets of queues, which are created automatically.

When the packets received from the network are mapped to queues, four access ingress-like queue-type and profile (rate-based) schedulers per destination card service the queues in strict priority. The following queue-type and profiled schedulers service the queues in the order listed:

1. Expedited in-profile scheduler
2. Best Effort in-profile scheduler
3. Expedited out-of-profile scheduler
4. Best Effort out-of-profile scheduler

To complete the operation, user-configurable shapers send the traffic into the fabric. See [Configurable ingress shaping to fabric \(access and network\)](#) for details. Throughout this operation, each packet retains its individual CoS value.

3.5.3.2 Network ingress 4-priority (Gen-3) scheduling

The adapter cards and ports that support 4-priority (Gen-3) scheduling for network ingress traffic are listed in [Table 12: Scheduling modes supported by adapter cards and ports at network ingress](#). See [QoS for Gen-3 adapter cards and platforms](#) for details.

3.5.3.3 Network ingress 16-priority scheduling

The cards and ports that support 16-priority scheduling for network ingress traffic are listed in [Table 12: Scheduling modes supported by adapter cards and ports at network ingress](#).

For a detailed description of how 16-priority scheduling functions, see [Network egress 16-priority scheduling](#).

The 7705 SAR-8 Shelf V2 and 7705 SAR-18 adapter cards, and the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-Wx ports provide sets of 8 queues for incoming traffic: 7 sets of queues for the 7705 SAR-8 Shelf V2, 17 sets of queues for the 7705 SAR-18, and 4 sets of queues for the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-Wx.

Each set of queues is specific to a destination adapter card. For the 7705 SAR-8 Shelf V2, 6 sets of queues are automatically created for each access egress adapter card, plus 1 set of queues for multicast traffic. For the 7705 SAR-18, 16 sets of queues are automatically created, plus 1 set of queues for multicast traffic. For the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-Wx, 3 sets of queues are automatically created, plus 1 set of queues for multicast traffic. For all these platforms, there is 1 additional set of queues for slow-path (control) traffic that is destined for the CSMs.

Each queue within each set provides buffer space for traffic isolation based on the classification carried out on EXP bits of the MPLS packet header (that is, the CoS setting).

All of the network ingress ports on an adapter card on a 7705 SAR-8 Shelf V2 or 7705 SAR-18 share the same sets of queues, which are created automatically. All of the network ingress ports across the entire 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, or 7705 SAR-Wx also share the same sets of queues, which are created automatically.

3.5.4 Network ingress shaping to fabric

After the traffic is scheduled, it must be sent to the fabric interface. To avoid congestion in the fabric and ease the effects of possible bursts, a shaper is implemented on each adapter card.

Network ingress shaping to the fabric operates in a similar fashion to access ingress shaping to the fabric. See [Ingress shaping to fabric \(access and network\)](#) for details.

3.5.5 Configurable network ingress shaping to fabric

Configuring network ingress shapers to the fabric is similar to configuring access ingress shapers to the fabric.

The ingress to-fabric shapers are user-configurable. For the 7705 SAR-8 Shelf V2 and the 7705 SAR-18, the maximum rate depends on a number of factors, including platform, chassis variant, and slot type. See [Configurable ingress shaping to fabric \(access and network\)](#) for details.

For information about fabric shapers on the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-Wx, see [Fabric shaping on the fixed platforms \(access and network\)](#). The 7705 SAR-X does not support configurable network ingress shapers.

3.5.6 Network fabric shaping on the fixed platforms

The 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, and 7705 SAR-Wx support user-configurable fabric shapers at rates of up to 5 Gb/s for access ingress traffic and network ingress traffic.

On the 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, and 7705 SAR-Wx, network ingress shapers to the fabric operate similarly to access ingress shapers to the fabric. The 7705 SAR-X does not support configurable network ingress shapers. See [Fabric shaping on the fixed platforms \(access and network\)](#) for more information.

3.6 Access egress

This section contains the following topics for traffic flow in the access egress direction:

- [Access egress queuing and scheduling](#)
- [Access egress per-SAP aggregate shapers \(access egress H-QoS\)](#)
- [Access egress shaping for hybrid ports](#)
- [Access egress for 4-priority \(Gen-3\) scheduling](#)
- [Access egress marking and re-marking](#)
- [Packet byte offset](#)

3.6.1 Access egress queuing and scheduling

The following sections discuss the queuing and scheduling of access egress traffic, which is traffic that egresses the fabric on the access side:

- [BMU traffic access egress queuing and scheduling](#)
- [ATM access egress queuing and scheduling](#)
- [Ethernet access egress queuing and scheduling](#)

Access egress scheduling takes place at the native traffic layer. As an example, when the ATM pseudowire payload is delivered from the network ingress to the access egress, the playback of the ATM cells to the appropriate ATM SAP is done according to ATM traffic management specifications.

Access egress scheduling is supported on the adapter cards and ports listed in the following table. The supported scheduling modes are 4-priority and 16-priority. The table shows which scheduling mode each card and port supports at access egress.



Note: For access ingress and egress, the 16-priority schedulers use additional hardware resources and capabilities, which results in increased throughput.

Table 13: Scheduling modes supported by adapter cards and ports at access egress

Adapter card or port	4-priority	16-priority
8-port Gigabit Ethernet Adapter card	✓	✓

Adapter card or port	4-priority	16-priority
Packet Microwave Adapter card	✓	✓
6-port Ethernet 10Gbps Adapter card ¹	✓	
10-port 1GigE/1-port 10GigE X-Adapter card (10-port 1GigE mode)	✓	✓
4-port SAR-H Fast Ethernet module	✓	
6-port SAR-M Ethernet module	✓	✓
Ethernet ports on the 7705 SAR-A	✓	✓
Ethernet ports on the 7705 SAR-Ax	✓	✓
Ethernet ports on the 7705 SAR-H	✓	✓
Ethernet ports on the 7705 SAR-Hc	✓	✓
Ethernet ports on the 7705 SAR-M	✓	✓
Ethernet ports on the 7705 SAR-Wx	✓	✓
Ethernet ports on the 7705 SAR-X ¹	✓	
16-port T1/E1 ASAP Adapter card	✓	
32-port T1/E1 ASAP Adapter card	✓	
2-port OC3/STM1 Channelized Adapter card	✓	
4-port OC3/STM1 Clear Channel Adapter card	✓	
4-port OC3/STM1 / 1-port OC12/STM4 Adapter card	✓	
4-port DS3/E3 Adapter card	✓	
T1/E1 ASAP ports on the 7705 SAR-A	✓	
T1/E1 ASAP ports on the 7705 SAR-M	✓	
TDM ports on the 7705 SAR-X	✓	
12-port Serial Data Interface card	✓	
6-port E&M Adapter card	✓	
6-port FXS Adapter card	✓	
8-port FXO Adapter card	✓	
8-port Voice & Teleprotection card	✓	
8-port C37.94 Teleprotection card	✓	

Adapter card or port	4-priority	16-priority
Integrated Services card	✓	

Note:

1. 4-priority scheduler for Gen-3 adapter card or platform.

3.6.1.1 BMU traffic access egress queuing and scheduling

At access egress, the 7705 SAR handles traffic management for unicast and BMU traffic in the same way. Unicast or BMU traffic is mapped to a queue and the mapping is based on the FC classification. Individual queues are then scheduled based on the available traffic.

3.6.1.2 ATM access egress queuing and scheduling

After the ATM pseudowire is terminated at the access egress, all the ATM cells are mapped to the default queue, which is queue 1, and queuing is performed per SAP. ATM access egress queuing and scheduling applies to the 16-port T1/E1 ASAP Adapter card, 32-port T1/E1 ASAP Adapter card, and 2-port OC3/STM1 Channelized Adapter card with atm/ima encapsulation. ATM access egress queuing and scheduling applies to the 4-port OC3/STM1 Clear Channel Adapter card and 4-port DS3/E3 Adapter card with atm encapsulation.

After the per-SAP queuing takes place, the ATM scheduler services these queues in the fashion and order defined below, based on the service categories assigned to each of these SAPs.

At access egress, CBR and rt-VBR VCs are always shaped because there is no option for the user to turn shaping off. Shaping for nrt-VBR is optional.

Strict priority scheduling in an exhaustive fashion takes place for the shaped VCs in the following order:

1. CBR (always shaped)
2. rt-VBR (always shaped)
3. nrt-VBR (when shaped, user-configurable for shaped or unshaped)

UBR traffic is not shaped. To offer maximum flexibility to the user, nrt-VBR unshaped (also known as scheduled) is implemented.

ATM traffic is serviced in priority order. CBR traffic has the highest priority and is serviced ahead of all other traffic. After all of the CBR traffic has been serviced, rt-VBR traffic is serviced. Then, nrt-VBR traffic is serviced.

After scheduling all the other traffic from the CBR and VBR service categories, UBR is serviced. If there is no other traffic, UBR can burst up to the line rate. Scheduled nrt-VBR is treated the same way as UBR. Both UBR and unshaped nrt-VBR are scheduled using the weighted round-robin scheduler.

The scheduler weight assigned to queues hosting scheduled nrt-VBR and UBR traffic is determined by the configured traffic rate. The weight used by the scheduler for UBR+ VCs is dependent on the minimum information rate (MIR) defined by the user. UBR with no MIR traffic has an MIR of 0.

Similarly, the scheduler weight is dependent on the sustained information rate (SIR) for scheduled nrt-VBR. Weight used by the scheduler is programmed automatically based on the user-configured MIR/SIR value and is not user-configurable.

For UBR+, the following tables are used to determine the weight of a UBR+ VC. These tables are also applicable to scheduled nrt-VBR weight determination. Instead of the MIR, the SIR is used to determine the scheduler weight.

Table 14: Scheduler weight values (WRR) based on MIR for T1/E1 ASAP Adapter cards and 2-port OC3/STM1 Channelized Adapter card

Minimum information rate	Scheduler weight
<64 kb/s	1
<128 kb/s	2
<256 kb/s	3
<512 kb/s	4
<1024 kb/s	5
<1536 kb/s	6
<1920 kb/s	7
≥1920 kb/s	8

Table 15: Scheduler weight values (WRR) based on MIR for the 4-port OC3/STM1 Clear Channel Adapter card

Range OC3 ATM	Range DS3 ATM	Weight
0 to 1 Mb/s	0 to 512 kb/s	1
>1 Mb/s to 4 Mb/s	>512 kb/s to 1 Mb/s	2
>4 Mb/s to 8 Mb/s	>1 Mb/s to 2 Mb/s	3
>8 Mb/s to 16 Mb/s	>2 Mb/s to 4 Mb/s	4
>16 Mb/s to 32 Mb/s	>4 Mb/s to 8 Mb/s	5
>32 Mb/s to 50 Mb/s	>8 Mb/s to 16 Mb/s	6
>50 Mb/s to 100 Mb/s	>16 Mb/s to 32 Mb/s	7
>100 Mb/s	>32 Mb/s	8

The access egress ATM scheduling behavior is shown in the following table. For UBR traffic, the scheduler weight of the lowest possible value is always used, which is the value of 1. Only cell-based operations are carried out.

Table 16: ATM scheduling and relative priorities

Flow type	Transmission rate	Priority
Shaped CBR	Limited to configured PIR	Strict priority over all other traffic
Shaped rt-VBR	Limited to configured SIR, but with bursts up to PIR within MBS	Strict priority over all but shaped CBR
Shaped nrt-VBR	Limited to configured SIR, but with bursts up to PIR within MBS	Strict priority over all scheduled traffic
Scheduled nrt-VBR	Weighted share (according to SIR) of port bandwidth remaining after shaped traffic has been exhausted	In the same WRR scheduler as UBR+ and UBR
Scheduled UBR+	Weighted share (according to MIR) of port bandwidth remaining after shaped traffic has been exhausted	In the same WRR scheduler as nrt-VBR and UBR
Scheduled UBR	Weighted share (with weight of 1) of port bandwidth remaining after shaped traffic has been exhausted	In the same WRR scheduler as nrt-VBR and UBR+

3.6.1.3 Ethernet access egress queuing and scheduling

Ethernet access egress queuing and scheduling is very similar to the Ethernet access ingress behavior. When the Ethernet pseudowire is terminated, traffic is mapped to up to eight different forwarding classes per SAP. Mapping traffic to different forwarding classes is performed based on the EXP bit settings of the received Ethernet pseudowire by network ingress classification.

Queue-type and profile scheduling are both supported for Ethernet access egress ports. If the queues are configured according to the tables and defaults described in this guide (implying a default mode of operation), the configuration is as follows:

- CoS-8 to CoS-5 Expedited in-profile
- CoS-4 to CoS-1 Best Effort in-profile
- CoS-8 to CoS-5 Expedited out-of-profile
- CoS-4 to CoS-1 Best Effort out-of-profile

In this default configuration, for queue-type scheduling, CoS-8 to CoS-5 are serviced by the Expedited scheduler, and CoS-4 to CoS-1 are serviced by the Best Effort scheduler. This default mode of operation can be altered to better fit the operating characteristics of specific SAPs.

With profile scheduling, the Ethernet frames can be either in-profile or out-of-profile, and scheduling takes into account the state of the Ethernet frames in conjunction with the configured CIR and PIR rates.

After the queuing, an aggregate queue-type and profile scheduling takes place in the following order:

1. Expedited in-profile traffic
2. Best Effort in-profile traffic
3. Expedited out-of-profile traffic

4. Best Effort out-of-profile traffic

After the traffic is scheduled using the aggregate queue-type and profile schedulers, the per-port shapers shape the traffic at a sub-rate (that is, at the configured/shaped port rate). Per-port shapers ensure that a sub-rate is met and attainable at all times.

3.6.2 Access egress per-SAP aggregate shapers (access egress H-QoS)

Per-SAP aggregate shapers in the access egress direction operate in a similar fashion to aggregate shapers for access ingress, except that egress traffic goes through the schedulers to the egress port shaper instead of through the schedulers to the fabric port as in the access ingress case. For information about how access egress and access ingress per-SAP shaping is similar, see [Access ingress per-SAP aggregate shapers \(access ingress H-QoS\)](#). For general information about per-SAP shapers, see [Per-SAP aggregate shapers \(H-QoS\) on Gen-2 hardware](#).

The arbitration of access egress traffic from the per-SAP aggregate shapers to the schedulers is described in the following section.

3.6.2.1 Access egress per-SAP shapers arbitration

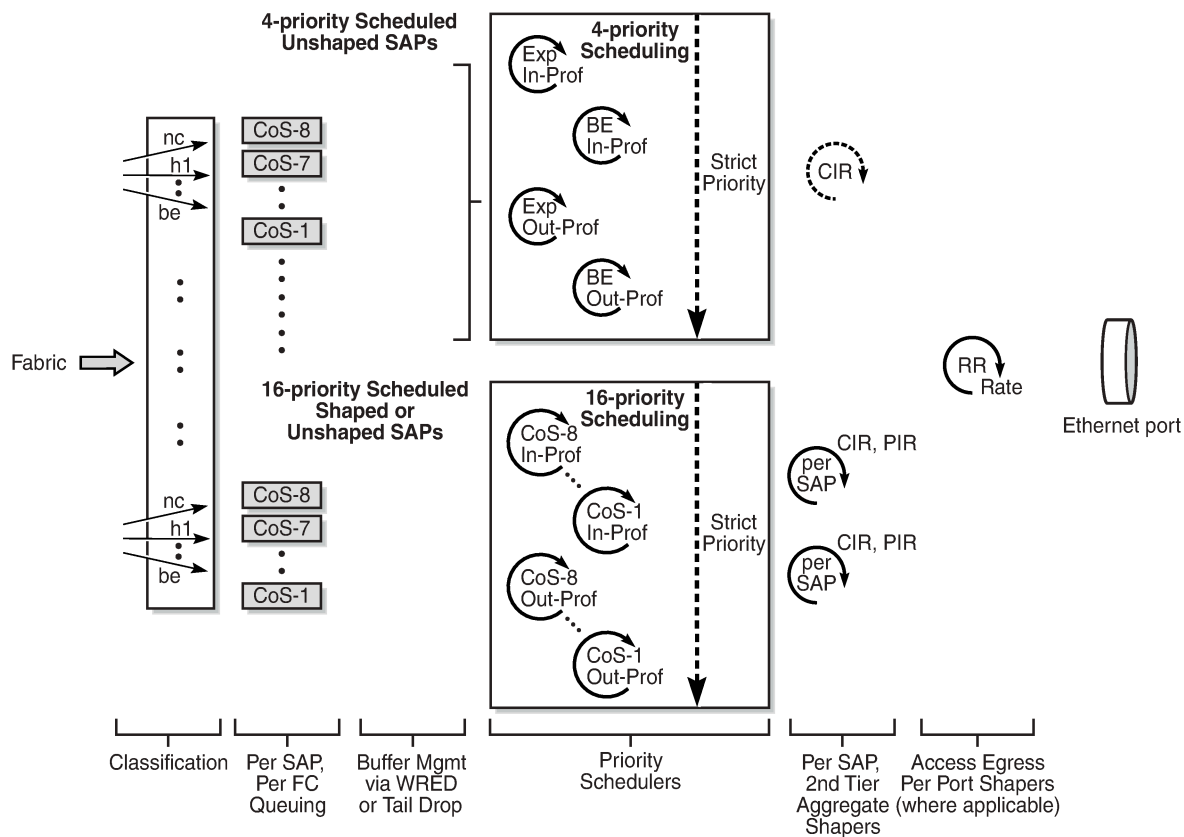
The arbitration of traffic from 4-priority and 16-priority schedulers toward an access egress port is achieved by configuring a committed aggregate rate limit for the aggregate of all the 4-priority unshaped SAPs. By configuring the 4-priority unshaped SAPs committed aggregate rate, the arbitration between the 16-priority shaped SAPs, 16-priority unshaped SAPs, and 4-priority unshaped SAPs is handled in the following priority order:

- committed traffic: 16-priority per-SAP **agg-rate-limit** committed for shaped SAPs and 4-priority aggregate committed rate for all the unshaped SAPs
- uncommitted traffic: 16-priority per-SAP **agg-rate-limit** uncommitted for shaped and unshaped SAPs and 4-priority aggregate uncommitted rate for all the unshaped SAPs

The following figure illustrates the traffic treatment for a single Ethernet port. It also illustrates that the shaped SAP aggregate CIR rate competes with the unshaped 4-priority aggregate CIR rate for port bandwidth. When the aggregate CIR rates are satisfied, the shaped SAP aggregate PIR rate competes with the 4-priority PIR rate (always maximum) for port bandwidth.

The egress aggregate CIR rate limit for all the unshaped 4-priority SAPs is configured using the **config>port>ethernet>access>egress>unshaped-sap-cir** command.

Figure 22: Access egress 16-priority and 4-priority per-SAP arbitration for a single port



23379

3.6.3 Access egress shaping for hybrid ports

Hybrid ports use a third-tier, dual-rate aggregate shaper to provide arbitration between the bulk of access and network egress traffic flows. For details, see [QoS for hybrid ports on Gen-2 hardware](#).

3.6.4 Access egress for 4-priority (Gen-3) scheduling

The adapter cards and ports that support 4-priority (Gen-3) scheduling for access egress traffic are listed in [Table 13: Scheduling modes supported by adapter cards and ports at access egress](#). See [QoS for Gen-3 adapter cards and platforms](#) for details.

3.6.5 Access egress marking and re-marking

At access egress, where the network-wide QoS boundary is reached, there may be a requirement to mark or re-mark the CoS indicators to match customer requirements. Dot1p and DSCP marking and re-marking is supported at Ethernet access egress.



Note: When dot1p re-marking is needed for a QinQ egress SAP, it may be necessary to use the **qinq-mark-top-only** command to indicate which qtag needs to have its dot1p bits re-marked. The **qinq-mark-top-only** command is found under the **config>service** context. See the 7705 SAR Services Guide, "VLL Services Command Reference", for details.

Similar to access ingress for Ethernet, DSCP marking or re-marking is supported for untagged, single-tagged, or double-tagged Ethernet frames.

On Ipipe SAPs over an Ethernet VLAN, both dot1p and DSCP marking and re-marking is supported at access egress. On Ipipe SAPs over PPP/MLPPP, DSCP marking and re-marking is supported at access egress. DSCP re-marking is supported for Ipipes using FR or cHDLC SAPS at the access egress.

3.6.6 Packet byte offset

Packet byte offset (PBO), or internal headerless rate, allows 7705 SAR schedulers to operate on a modified packet size by adding or subtracting a certain number of bytes. The actual packet size remains the same but schedulers take into account the modified size as opposed to the actual size of the packet. One of the main uses of the packet byte offset feature is to allow scheduling, at access ingress, to be carried out on the received packet size without taking into account service (for example, PW, MPLS) or internal overhead. Transport providers who sell bandwidth to customers typically need the 7705 SAR shapers/schedulers to only take into account the received packet size without the added overhead in order to accurately calculate the bandwidth they need to provide to their customers. Packet byte offset addresses this requirement. Another common use is at egress where port shapers can take into account four additional bytes, associated with Ethernet FCS.

Packet byte offset is configured under QoS profiles. Packet size modification may be desired to accommodate inclusion or exclusion of certain headers or even fields of headers during the scheduling operation. The packet size that the schedulers take into account is altered to accommodate or omit the desired number of bytes. Both addition and subtraction options are supported by the **packet-byte-offset** command. The actual packet size is not modified by the command; only the size used by ingress or egress schedulers is changed. The scheduling rates are affected by the offset, as well as the statistics (accounting) associated with the queue. Packet byte offset does not affect port-level and service-level statistics. It only affects the queue statistics.

When a QoS policy configured with packet byte offset is applied to a SAP or network interface, all the octet counters and statistics operate and report based on the new adjusted value. If configured, per-SAP aggregate shapers and per-customer aggregate shapers also operate on the adjusted packet sizes. The only exceptions to this rule are port shapers. The egress port shapers do not take the adjusted packet size into account but operate only on the final packet size.



Note: The fabric shaper, in general, does not take the adjusted packet size into account.

The following table shows PBO support on the 7705 SAR.

Table 17: PBO for SAPs and platforms

Traffic direction and PBO count				Second and third generation adapter cards and platforms
	Per SAP CoS Queue	Per SAP Shaper	Per Customer Shaper	Fabric Shaper 7705 SAR-8 Shelf V2 / 7705 SAR-18

Traffic direction and PBO count				Second and third generation adapter cards and platforms	
				Sum of adjusted MSS shapers \leq fabric shapers	Sum of adjusted MSS shapers > fabric shapers
Access ingress	✓	✓	✓	internal packet size, no FCS	internal packet size, no FCS
auto	3	3	3	3	internal packet size, no FCS (fabric shaper rate)
add 50	2	2	2	2	3
subtract 50	6	6	6	6	3
	Per CoS Queue	Bypass	Fabric Shaper (on Non-Chassis Based Nodes) Otherwise Bypass	Fabric Shaper 7705 SAR-8 Shelf V2 / 7705 SAR-18	
Network ingress	✓	n/a	✓	✓	
add 50	2	n/a	2	2	
subtract 50	6	n/a	6	6	
	Per SAP CoS Queue	Per SAP Shaper	Per Customer Shaper	Port Shaper	
				Sum of adjusted MSS shapers \leq egress rate	Sum of adjusted MSS shapers > egress rate
Access egress	✓	✓	✓	✓	final packet size, FCS optional
add 50	2	2	2	2 (room for 3)	3
subtract 50	6	6	6	6	3
	Per CoS Queue	Per VLAN Shaper	Bypass	Port Shaper	

Traffic direction and PBO count				Second and third generation adapter cards and platforms	
				Sum of adjusted VLAN shapers ≤ egress rate	Sum of adjusted VLAN shapers > egress rate
Network egress	✓	✓	n/a	✓	final packet size, FCS optional
add 50	2	2	n/a	2 (room for 3)	2 (room for 3)
subtract 50	6	6	n/a	6	3
	Per Access / Network CoS queue	SAP / VLAN Shaper	Per Customer Shaper Access / Network Arbitrator	Port Shaper	
				Sum of adjusted MSS/NW arbitrator shapers ≤ egress rate	Sum of adjusted MSS/NW arbitrator shapers > egress rate
Hybrid egress	✓	✓	✓	✓	final packet size, FCS optional
add 50	2	2	2	2 (room for 3)	3
subtract 50	6	6	6	6	3

3.7 QoS policies overview

This section contains the following topics related to QoS policies:

- [Overview](#)
- [Service ingress QoS policies](#)
- [Service egress QoS policies](#)
- [MC-MLPPP SAP egress QoS policies](#)
- [Network and network queue QoS policies](#)
- [Network and service QoS queue parameters](#)
- [Slope policies \(WRED and RED\)](#)
- [ATM traffic descriptor profiles](#)
- [Fabric profiles](#)

- [Shaper policies](#)
- [QoS policy entities](#)

3.7.1 Overview

7705 SAR QoS policies are applied on service ingress, service egress, and network interfaces. The service ingress and service egress points may be considered as the network QoS boundaries for the service being provided.

The QoS policies define:

- classification rules for how traffic is mapped to forwarding classes
- how forwarding classes are aggregated under queues
- the queue parameters used for policing, shaping, and buffer allocation
- QoS marking/interpretation

There are several types of QoS policies (see [Table 18: QoS policy types and descriptions](#) for summaries and references to details):

- service ingress (also known as access ingress)
- service egress (also known as access egress)
- MC-MLPPP SAP egress
- network (for ingress and egress and ring)
 - IP interface type policy for network ingress and egress
 - ring type policy for Ethernet bridging domain on a ring adapter card
- network queue (for ingress and egress)
- slope
- ATM traffic descriptor profile
- fabric profile
- shaper



Note: The terms access ingress/egress and service ingress/egress are interchangeable. The previous sections used the term access, and the sections that follow use the term service.

Service ingress QoS policies are applied to the customer-facing SAPs and map traffic to forwarding class queues on ingress. The mapping of traffic to queues can be based on combinations of customer QoS marking (dot1p bits and DSCP values). The number of forwarding class queues for ingress traffic and the queue characteristics are defined within the policy. There can be up to eight ingress forwarding class queues in the policy, one for each forwarding class.

Within a service ingress QoS policy, up to three queues per forwarding class can be used for multipoint traffic for multipoint services. Multipoint traffic consists of broadcast, multicast, and unknown (BMU) traffic types. For VPLS, four types of forwarding are supported (which are not to be confused with forwarding classes): unicast, broadcast, multicast, and unknown. The BMU types are flooded to all destinations within the service, while the unicast forwarding type is handled in a point-to-point fashion within the service.

Service ingress QoS policies on the 7705 SAR allow flexible arrangement of these queues. For example, more than one FC can be mapped to a single queue, both unicast and multipoint (BMU) traffic can be

mapped to a single queue, or unicast and BMU traffic can be mapped to separate queues. Therefore, customers are not limited to the default configurations that are described in this guide.

Service egress QoS policies are applied to egress SAPs and provide the configurations needed to map forwarding classes to service egress queues. Each service can have up to eight queues configured, since a service may require multiple forwarding classes. A service egress QoS policy also defines how to re-mark dot1p bits and DSCP values of the customer traffic in native format based on the forwarding class of the customer traffic.

Network ingress and egress QoS policies are applied to network interfaces. On ingress for traffic received from the network, the policy maps incoming EXP values to forwarding classes and profile states. On egress, the policy maps forwarding classes and profile states to EXP values for traffic to be transmitted into the network.

On the network side, there are two types of QoS policies: network and network queue (see [Table 18: QoS policy types and descriptions](#)). The network type of QoS policy is applied to the network interface under the **config>router>interface** command and contains the EXP marking rules for both ingress and egress. The network queue type of QoS policy defines all of the internal settings; that is, how the queues, or sets of queues (for ingress), are set up and used per physical port on egress and per adapter card for ingress.

A ring type network policy can be applied to the ring ports and the add/drop port on a ring adapter card. The policy is created under the **config>qos>network** command, and applied at the adapter card level under the **config>card>mda** command. The policy maps each dot1p value to a queue and a profile state.

If GRE or IP tunneling is enabled, policy mapping can be set up to use DSCP bits.

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

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

One service ingress QoS policy and one service egress QoS policy can be applied to a specific SAP. One network QoS policy can be applied to a specific interface. A network QoS policy defines both ingress and egress behavior. If no QoS policy is explicitly applied to a SAP or network interface, a default QoS policy is applied.

The following table provides a summary of the major functions performed by the QoS policies.

Table 18: QoS policy types and descriptions

Policy type	Applied at...	Description	Section
Service Ingress	SAP ingress	Defines up to eight forwarding class queues and queue parameters for traffic classification Defines match criteria to map flows to the queues based on combinations of customer QoS (dot1p bits and DSCP values)	Service ingress QoS policies
Service Egress	SAP egress	Defines up to eight forwarding class queues and queue parameters for traffic classification Maps one or more forwarding classes to the queues	Service egress QoS policies

Policy type	Applied at...	Description	Section
MC-MLPPP	SAP egress	Defines up to eight forwarding class queues and queue parameters for traffic classification Maps one or more forwarding classes to the queues	MC-MLPPP SAP egress QoS policies
Network	Network interface	Packets are marked using QoS policies on edge devices, such as the 7705 SAR at access ingress. Invoking a QoS policy on a network port allows for the packets that match the policy criteria to be re-marked at network egress for appropriate CoS handling across the network	Network QoS policies
Network Queue	Adapter card network ingress and egress	Defines forwarding class mappings to network queues	Network queue QoS policies
Slope	Adapter card ports	Enables or disables the high-slope and low-slope parameters within the egress or ingress queue	Slope policies (WRED and RED)
ATM Traffic Descriptor Profile	SAP ingress	Defines the expected rates and characteristics of traffic. Specified traffic parameters are used for policing ATM cells and for selecting the service category for the per-VC queue.	ATM traffic descriptor profiles
	SAP egress	Defines the expected rates and characteristics of traffic. Specified traffic parameters are used for scheduling and shaping ATM cells and for selecting the service category for the per-VC queue.	ATM traffic descriptor profiles
Fabric Profile	Adapter card access and network ingress	Defines access and network ingress to-fabric shapers at user-configurable rates	Fabric profiles
Shaper	Adapter card ports	Defines dual-rate shaping parameters for a shaper group in a shaper policy	Shaper policies

3.7.2 Service ingress QoS policies

Service ingress QoS policies define ingress service forwarding class queues and map flows to those queues. When a service ingress QoS policy is created, it always has a default ingress traffic queue defined that cannot be deleted. These queues exist within the definition of the policy. The queues only get created when the policy is applied to a SAP.

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

- a unique service ingress QoS policy ID
- a QoS policy scope of template or exclusive

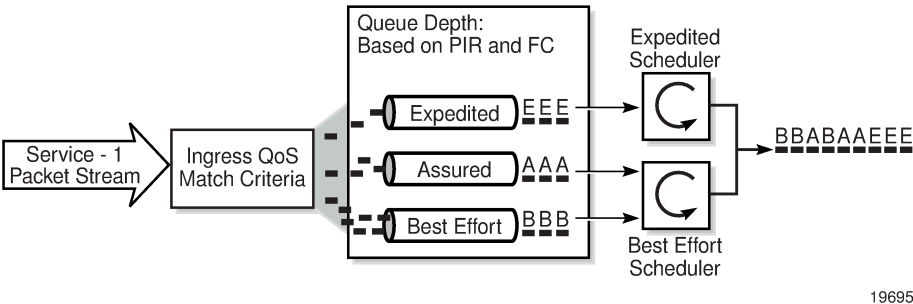
- at least one default ingress forwarding class queue. The parameters that can be configured for a queue are discussed in [Network and service QoS queue parameters](#).

Optional service ingress QoS policy elements include:

- additional ingress queues up to a total of eight
- QoS policy match criteria to map packets to a forwarding class

Each queue can have unique queue parameters to allow individual policing and rate shaping of the flow mapped to the forwarding class. The following figure depicts service traffic being classified into three different forwarding class queues.

Figure 23: Traffic queuing model for three queues and three classes



The mapping of flows to forwarding classes is controlled by comparing each packet to the match criteria in the QoS policy. The ingress packet classification to forwarding class and enqueueing priority is subject to a classification hierarchy. Each type of classification rule is interpreted with a specific priority in the hierarchy. The following table is an example for an Ethernet SAP (that is, a SAP defined over a whole Ethernet port, over a single VLAN, or over QinQ VLANs). The table lists the classification rules in the order in which they are evaluated.

Table 19: Forwarding class and enqueueing priority classification hierarchy based on rule type

Rule	Forwarding class	Enqueueing priority	Comments
default-fc	Set to the policy's default FC	Set to the policy default	All packets match the default rule
dot1p dot1p-value	Set when an <i>fc-name</i> exists in the policy Otherwise, preserve from the previous match	Set when the <i>priority</i> parameter is high or low Otherwise, preserve from the previous match	Each <i>dot1p-value</i> must be explicitly defined. Each packet can only match a single dot1p rule. For QinQ applications, the <i>dot1p-value</i> used (top or bottom) is specified by the match-qinq-dot1p command.
dscp dscp-name	Set when an <i>fc-name</i> exists in the policy Otherwise, preserve from the previous match	Set when the <i>priority</i> parameter is high or low in the entry Otherwise, preserve from the previous match	Each <i>dscp-name</i> that defines the DSCP value must be explicitly defined. Each packet can only match a single DSCP rule.

The enqueueing priority is specified as part of the classification rule and is set to high or low. The enqueueing priority relates to the forwarding class queue's high-priority-only allocation, where only packets with a high enqueueing priority are accepted into the queue when the queue's depth reaches the defined threshold. See [High-priority-only buffers](#).

The mapping of ingress traffic to a forwarding class based on dot1p or DSCP bits is optional. The default service ingress policy is implicitly applied to all SAPs that do not explicitly have another service ingress policy assigned. The characteristics of the default policy are listed in the following table.

Table 20: Default service ingress policy ID 1 definition

Characteristic	Item	Definition
Queues	Queue 1	One queue for all ingress traffic: <ul style="list-style-type: none"> Forwarding Class: Best Effort (BE) CIR = 0 PIR = max (line rate) MBS = default (180 kB) CBS = default (8 kB for 512 byte buffer size, 18 kB for 2304 byte buffer size) ¹ HP Only = default (10%)
Flows	Default FC	One flow defined for all traffic: <ul style="list-style-type: none"> all traffic mapped to Best Effort (BE) with a low priority

Note:

- See [Table 4: Buffer support on adapter cards and platforms](#) for a list of adapter cards and buffer sizes.

3.7.3 Service egress QoS policies

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

- per-service egress shaping, soft-policing capabilities
- more granular, more fair scheduling per service into the access network
- per-service statistics for forwarded and discarded service packets

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

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

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

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

- a unique service egress QoS policy ID
- a QoS policy scope of template or exclusive
- at least one defined default queue. The parameters that can be configured for a queue are discussed in [Network and service QoS queue parameters](#).

Optional service egress QoS policy elements include:

- additional queues, up to a total of eight separate queues
- dot1p priority and DSCP value re-marking based on forwarding class

Each queue in a policy is associated with one or more of the supported forwarding classes. Each queue can have its individual queue parameters, allowing individual rate shaping of the forwarding classes mapped to the queue. More complex service queuing models are supported in the 7705 SAR where each forwarding class is associated with a dedicated queue.

The forwarding class determination per service egress packet is determined at ingress. If the packet ingressed the service on the same 7705 SAR router, the service ingress classification rules determine the forwarding class of the packet. If the packet was received over a service transport tunnel, the forwarding class is marked in the tunnel transport encapsulation.

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

The default service egress policy is applied to all SAPs that do not have another service egress policy explicitly assigned. The characteristics of the default policy are listed in the following table.

Table 21: Default service egress policy ID 1 definition

Characteristic	Item	Definition
Queues	Queue 1	One queue defined for all traffic classes: <ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • MBS = default (180 kB) • CBS = default (8 kB for 512 byte buffer size, 18 kB for 2304 byte buffer size) ¹ • HP Only = default (10%)
Flows	Default action	One flow defined for all traffic classes: <ul style="list-style-type: none"> • all traffic mapped to queue 1 with no marking of IEEE 802.1p or DSCP values

Note:

1. See [Table 4: Buffer support on adapter cards and platforms](#) for a list of adapter cards and buffer sizes.

3.7.4 MC-MLPPP SAP egress QoS policies

SAPs running MC-MLPPP have their own SAP egress QoS policies that differ from standard policies. Unlike standard SAP policies, MC-MLPPP SAP egress policies do not contain queue types, CIR, CIR adaptation rules, or dot1p re-marking.

Standard and MC-MLPPP SAP egress policies can never have the same policy ID except when the policy ID is 1 (default). Standard SAP egress QoS policies cannot be applied to SAPs running MC-MLPPP. Similarly, MC-MLPPP SAP egress QoS policies cannot be applied to standard SAPs. The default policy can be applied to both MC-MLPPP and other SAPs. It will remain the default policy regardless of SAP type.

MC-MLPPP on the 7705 SAR supports scheduling based on multiclass implementation. Instead of the standard profiled queue-type scheduling, an MC-MLPPP encapsulated access port performs class-based traffic servicing.

The four MC-MLPPP classes are scheduled in a strict priority fashion, as shown in the following table.

Table 22: MC-MLPPP class priorities

MC-MLPPP class	Priority
0	Priority over all other classes
1	Priority over classes 2 and 3
2	Priority over class 3
3	No priority

For example, if a packet is sent to an MC-MLPPP class 3 queue and all other queues are empty, the 7705 SAR fragments the packet according to the configured fragment size and begins sending the fragments. If a new packet is sent to an MC-MLPPP class 2 queue, the 7705 SAR finishes sending any fragments of the class 3 packet that are on the wire, then holds back the remaining fragments in order to service the higher-priority packet. The fragments of the first packet remain at the top of the class 3 queue. For packets of the same class, MC-MLPPP class queues operate on a first-in, first-out basis.

The user configures the required number of MLPPP classes to use on a bundle. The forwarding class of the packet, as determined by the ingress QoS classification, is used to determine the MLPPP class for the packet. The mapping of forwarding class to MLPPP class is a function of the user-configurable number of MLPPP classes. The default mapping for a 4-class, 3-class, and 2-class MLPPP bundle is shown in the following table.

Table 23: Packet forwarding class to MLPPP class mapping

FC ID	FC name	MLPPP class 4-class bundle	MLPPP class 3-class bundle	MLPPP class 2-class bundle
7	NC	0	0	0
6	H1	0	0	0
5	EF	1	1	1

FC ID	FC name	MLPPP class 4-class bundle	MLPPP class 3-class bundle	MLPPP class 2-class bundle
4	H2	1	1	1
3	L1	2	2	1
2	AF	2	2	1
1	L2	3	2	1
0	BE	3	2	1

If one or more forwarding classes are mapped to a queue, the scheduling priority of the queue is based on the lowest forwarding class mapped to it. For example, if forwarding classes 0 and 7 are mapped to a queue, the queue is serviced by MC-MLPPP class 3 in a 4-class bundle model.

3.7.5 Network and network queue QoS policies

The QoS mechanisms within the 7705 SAR are specialized for the type of traffic on the interface. For customer interfaces, there is service ingress and service egress traffic, and for network interfaces, there is network ingress and network egress traffic.

The 7705 SAR uses QoS policies applied to a SAP for a service or to a network port to define the queuing, queue attributes, and QoS marking/interpretation.

The 7705 SAR supports the following types of network and service QoS policies:

- [Network QoS policies](#)
- [Network queue QoS policies](#)
- [Service ingress QoS policies](#) (described previously)
- [Service egress QoS policies](#) (described previously)



Note: Queuing parameters are the same for both network and service QoS policies. See [Network and service QoS queue parameters](#).

3.7.5.1 Network QoS policies

Network QoS policies define egress QoS marking and ingress QoS classification for traffic on network interfaces. The 7705 SAR automatically creates egress queues for each of the forwarding classes on network interfaces.

A network QoS policy defines ingress, egress, and ring handling of QoS on the network interface. The following functions are defined:

- ingress
 - defines label switched path Experimental bit (LSP EXP) value mappings to forwarding classes
 - defines DSCP name mappings to forwarding classes
- egress

- defines forwarding class to LSP EXP and dot1p value markings
- defines forwarding class to DSCP value markings
- ring
 - defines dot1p bit value mappings to queue and profile state

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

- a unique network QoS policy ID
- egress forwarding class to LSP EXP value mappings for each forwarding class used
- a default ingress forwarding class and in-profile/out-of-profile state
- a default queue and in-profile/out-of-profile state for ring type network QoS policy

Optional ip-interface type network QoS policy elements include the LSP EXP value or DSCP name to forwarding class and profile state mappings for all EXP values or DSCP values received. Optional ring type network QoS policy elements include the dot1p bits value to queue and profile state mappings for all dot1p bit values received.

Network policy ID 1 is reserved as the default network QoS policy. The default policy cannot be deleted or changed. The default network QoS policy is applied to all network interfaces and ring ports (for ring adapter cards) that do not have another network QoS policy explicitly assigned.

The following tables list the various network QoS policy default mappings:

- [Table 24: Default network QoS policy egress marking](#)
- [Table 25: Default network QoS policy DSCP-to-forwarding class mappings](#)
- [Table 26: Default network QoS policy LSP EXP-to-forwarding class mappings](#)
- [Table 27: Default network QoS policy dot1p-to-queue class mappings](#)

The following table lists the default mapping of forwarding class to LSP EXP values and DSCP names for network egress.

Table 24: Default network QoS policy egress marking

FC-ID	FC name	FC label	DiffServ name	Egress LSP EXP marking		Egress DSCP marking	
				In-profile	Out-of-profile	In-profile name	Out-of-profile name
7	Network Control	nc	NC2	111 - 7	111 - 7	nc2 111000 - 56	nc2 111000 - 56
6	High-1	h1	NC1	110 - 6	110 - 6	nc1 110000 - 48	nc1 110000 - 48
5	Expedited	ef	EF	101 - 5	101 - 5	ef 101110 - 46	ef 101110 - 46
4	High-2	h2	AF4	100 - 4	100 - 4	af41 100010 - 34	af42 100100 - 36

FC-ID	FC name	FC label	DiffServ name	Egress LSP EXP marking		Egress DSCP marking	
				In-profile	Out-of-profile	In-profile name	Out-of-profile name
3	Low-1	l1	AF2	011 - 3	010 - 2	af21 010010 - 18	af22 010100 - 20
2	Assured	af	AF1	011 - 3	010 - 2	af11 001010 - 10	af12 001100 - 12
1	Low-2	l2	CS1	001 - 1	001 - 1	cs1 001000 - 8	cs1 001000 - 8
0	Best Effort	be	BE	000 - 0	000 - 0	be 000000 - 0	be 000000 - 0



Note: IP traffic from network interfaces is always trusted. There is no re-marking performed on network egress for global routing table (GRT) forwarded IP traffic (with the exception of GRE and IPSec tunnels). This also applies to IES interface traffic that is forwarded to the network egress interface.

For network ingress, the following table lists the default mapping of DSCP name to forwarding class and profile state for the default network QoS policy.

Table 25: Default network QoS policy DSCP-to-forwarding class mappings

Ingress DSCP			Forwarding class		
DSCP name	DSCP value	FC ID	Name	Label	Profile state
Default ¹		0	Best-Effort	be	Out
ef	101110 - 46	5	Expedited	ef	In
cs1	001000 - 8	1	Low-2	l2	In
nc-1	110000 - 48	6	High-1	h1	In
nc-2	111000 - 56	7	Network Control	nc	In
af11	001010 - 10	2	Assured	af	In
af12	001100 - 12	2	Assured	af	Out
af13	001110 - 14	2	Assured	af	Out
af21	010010 - 18	3	Low-1	l1	In
af22	010100 - 20	3	Low-1	l1	Out

Ingress DSCP			Forwarding class		
DSCP name	DSCP value	FC ID	Name	Label	Profile state
af23	010110 - 22	3	Low-1	l1	Out
af31	011010 - 26	3	Low-1	l1	In
af32	011100 - 28	3	Low-1	l1	Out
af33	011110 - 30	3	Low-1	l1	Out
af41	100010 - 34	4	High-2	h2	In
af42	100100 - 36	4	High-2	h2	Out
af43	100110 - 38	4	High-2	h2	Out

Note:

1. The default forwarding class mapping is used for all DSCP name values for which there is no explicit forwarding class mapping.

The following table lists the default mapping of LSP EXP values to forwarding class and profile state for the default network QoS policy.

Table 26: Default network QoS policy LSP EXP-to-forwarding class mappings

Ingress LSP EXP			Forwarding class		
LSP EXP ID	LSP EXP value	FC ID	Name	Label	Profile state
Default ¹		0	Best-Effort	be	Out
1	001 - 1	1	Low-2	l2	In
2	010 - 2	2	Assured	af	Out
3	011 - 3	2	Assured	af	In
4	100 - 4	4	High-2	h2	In
5	101 - 5	5	Expedited	ef	In
6	110 - 6	6	High-1	h1	In
7	111 - 7	7	Network Control	nc	In

Note:

1. The default forwarding class mapping is used for all LSP EXP values for which there is no explicit forwarding class mapping.

The following table lists the default mapping of dot1p values to queue and profile state for the default network QoS policy.

Table 27: Default network QoS policy dot1p-to-queue class mappings

Dot1p value	Queue	Profile state
0	1	Out
1	2	In
2	3	Out
3	3 ¹	In
4	5	In
5	6	In
6	7	In
7	8	In

Note:

1. The default queue mapping for dot1p values 2 and 3 are both queue 3.

3.7.5.1.1 CoS marking for self-generated traffic

The 7705 SAR is the source of some types of traffic; for example, a link state PDU for sending IS-IS topology updates or an SNMP trap sent to indicate that an event has happened. This type of traffic that is created by the 7705 SAR is considered to be self-generated traffic (SGT). Another example of self-generated traffic is Telnet, but in that application, user commands initiate the sending of the Telnet traffic.

Network operators often have different QoS models throughout their networks and apply different QoS schemes to portions of the networks to better accommodate delay, jitter, and loss requirements of different applications. The class of service (DSCP or dot1p) bits of self-generated traffic can be marked on a per-application basis to match the network operator's QoS scheme. This marking option enhances the ability of the 7705 SAR to match the various requirements of these applications.

The 7705 SAR supports marking self-generated traffic for the base routers and for virtual routers. See "QoS Policies" in the 7705 SAR Services Guide for information about SGT QoS as applied to virtual routers (for VPRN services).

The DSCP and dot1p values of the self-generated traffic, where applicable, are marked in accordance with the values that are configured under the **sgt-qos** command. In the egress direction, self-generated traffic is forwarded using the egress control queue to ensure premium treatment, unless SGT redirection is configured (see [SGT redirection](#)). PTP (IEEE 1588v2) and SAA-enabled ICMP traffic is forwarded using the CoS queue. The next-hop router uses the DSCP values to classify the traffic accordingly.



Note: IS-IS and ARP traffic are not IP-generated traffic types and are not DSCP-configurable; however, the dot1p bits can be configured in the same way as the DSCP bits. The default setting for the dot1p bits for both types of traffic is 111. For all other applications, the dot1p bits are marked based on the mapped network egress forwarding class.

The following table lists various applications and indicates whether they have configurable DSCP or dot1p markings.

Table 28: Applications and support for configurable DSCP or dot1p markings

Application	Supported marking	Default DSCP/dot1p
ARP	dot1p	7
IS-IS	dot1p	7
BGP	DSCP	NC1
DHCP	DSCP	NC1
DNS	DSCP	AF41
FTP	DSCP	AF41
ICMP (ping)	DSCP	BE
IGMP	DSCP	NC1
LDP (T-LDP)	DSCP	NC1
MCFW	DSCP	NC1
MLD	DSCP	NC1
NDIS	DSCP	NC1
NTP	DSCP	NC1
OSPF	DSCP	NC1
PIM	DSCP	NC1
1588 PTP	DSCP	NC1
RADIUS	DSCP	AF41
RIP	DSCP	NC1
RSVP	DSCP	NC1
SNMP (get, set, etc.)	DSCP	AF41
SNMP trap/log	DSCP	AF41
SSH (SCP)	DSCP	AF41
syslog	DSCP	AF41
TACACS+	DSCP	AF41
Telnet	DSCP	AF41

Application	Supported marking	Default DSCP/dot1p
TFTP	DSCP	AF41
Traceroute	DSCP	BE
VRRP	DSCP	NC1

**Note:**

- PTP in the context of SGT QoS is defined as Precision Timing Protocol and is an application in the 7705 SAR. The PTP application name is also used in areas such as event-control and logging. Precision Timing Protocol is defined in IEEE 1588-2008.
- PTP in the context of IP filters is defined as Performance Transparency Protocol. IP protocols can be used as IP filter match criteria; the match is made on the 8-bit protocol field in the IP header.

3.7.5.1.2 SGT redirection

The 7705 SAR can be used in deployments where the uplink bandwidth capacity is considerably less than if the router is used for fixed or mobile backhaul applications. However, the 7705 SAR is optimized to operate in environments with megabits per second of uplink capacity for network operations. Therefore, many of the software timers are designed to ensure the fastest possible detection of failures, without considering bandwidth limitations. In deployments with very low bandwidth constraints, the system must also be optimized for effective operation of the routers without any interruption to mission-critical customer traffic.

In lower-bandwidth deployments, SGT can impact mission-critical user traffic such as TDM pseudowire traffic. To minimize the impact on this traffic, SGT can be redirected to a data queue rather than to the high-priority control queue on egress. All SGT applications can be redirected to a data queue, but the type of application must be considered because not all SGT is suitable to be scheduled at a lower priority. SGT applications such as FTP, TFTP, and syslog can be mapped to a lower-priority queue.



Caution: Care must be taken when determining which SGT applications should be moved to data queues, as interrupting traffic flow for applications such as routing protocols (for example, BGP, OSPF, and IS-IS) and MPLS can adversely affect router operation, services, and the network.

As an example, in a scenario where the uplink bandwidth is limited to a fractional E1 link with 2 x DS0 channel groups, downloading software for a new release can disrupt TDM pseudowire traffic, especially if SGT traffic is always serviced first over all other traffic flows. Having the option to map a subset of SGT to data queues will ensure that the mission-critical traffic flows effectively. For example, if FTP traffic is redirected to the best-effort forwarding queue, FTP traffic is then serviced only after all higher-priority traffic is serviced, including network control traffic and TDM pseudowire traffic. This redirection ensures the correct treatment of all traffic types matching the requirements of the network.



Caution: Timeouts for signaling and/or routing protocols can initiate a session teardown. The teardown will not only have local impacts as severe as losing the sole uplink of a node but could also be network-wide if loss of a node is propagated throughout the network. In this scenario, routing protocols, as an example, will rerun on all the nodes within the same area, generating a CPU load and extra control-plane traffic. Such a scenario could cause potential instability across the network or an area of the network. To avoid this scenario, the user must ensure that uplink

capacity is enough to transmit all crucial SGT plus the mission-critical user traffic, and SGT redirection should only be used for the non-mission-critical traffic that can tolerate delay and jitter.

Redirection of SGT applications is done using the **config>router>sgt-qos>application>fc-queue** or **config>service>vprn>sgt-qos>application>fc-queue** command.

Redirection of the global ping application is not done through the **sgt-qos** menu hierarchy; this is configured using the **fc-queue** option in the **ping** command. See the 7705 SAR OAM and Diagnostics Guide, "OAM and SAA Command Reference", for details.

SGT redirection is supported on the base router and the virtual routers on ports with Ethernet or PPP/MLPPP encapsulation.

3.7.5.2 Network queue QoS policies

Network queue policies define the queue characteristics that are used in determining the scheduling and queuing behavior for a forwarding class. Network queue policies are applied on ingress and egress network ports as well as on the ring ports and the add/drop port on the 2-port 10GigE (Ethernet) Adapter card and 2-port 10GigE (Ethernet) module.

Network queue policies are identified with a unique policy name that conforms to the standard 7705 SAR alphanumeric naming conventions. The policy name is user-configured when the policy is created.

Network queue policies can be configured to use up to 16 queues (8 unicast and 8 multicast). This means that the number of queues can vary. Not all user-created policies will require and use 16 queues; however, the system default network queue policy (named "default") does define 16 queues.

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

- committed buffer size (CBS) as a percentage of the buffer pool
- maximum buffer size (MBS) as a percentage of the buffer pool
- high-priority-only buffers as a percentage of MBS
- peak information rate (PIR) as a percentage of egress port bandwidth
- committed information rate (CIR) as a percentage of egress port bandwidth

The following table describes the default network queue policy definition.



Note:

- The system default network queue policy cannot be modified or deleted.
- In the table, the value for Rate in the Definition column is the PIR value.

Table 29: Default network queue policy definition

Forwarding class	Queue	Definition	Queue	Definition
Network-Control (nc)	8	Rate = 100% CIR = 10% MBS = 2.5% CBS = 0.25% High-Prio-Only = 10%	16	Rate = 100% CIR = 10% MBS = 2.5% CBS = 0.1% High-Prio-Only = 10%

Forwarding class	Queue	Definition	Queue	Definition
High-1 (h1)	7	Rate = 100% CIR = 10% MBS = 2.5% CBS = 0.25% High-Prio-Only = 10%	15	Rate = 100% CIR = 10% MBS = 2.5% CBS = 0.1% High-Prio-Only = 10%
Expedited (ef)	6	Rate = 100% CIR = 100% MBS = 5% CBS = 0.75% High-Prio-Only = 10%	14	Rate = 100% CIR = 100% MBS = 5% CBS = 0.1% High-Prio-Only = 10%
High-2 (h2)	5	Rate = 100% CIR = 100% MBS = 5% CBS = 0.75% High-Prio-Only = 10%	13	Rate = 100% CIR = 100% MBS = 5% CBS = 0.1% High-Prio-Only = 10%
Low-1 (l1)	4	Rate = 100% CIR = 25% MBS = 2.5% CBS = 0.25% High-Prio-Only = 10%	12	Rate = 100% CIR = 5% MBS = 2.5% CBS = 0.25% High-Prio-Only = 10%
Assured (af)	3	Rate = 100% CIR = 25% MBS = 5% CBS = 0.75% High-Prio-Only = 10%	11	Rate = 100% CIR = 5% MBS = 5% CBS = 0.1% High-Prio-Only = 10%
Low-2 (l2)	2	Rate = 100% CIR = 25% MBS = 5% CBS = 0.25% High-Prio-Only = 10%	10	Rate = 100% CIR = 5% MBS = 5% CBS = 0.1% High-Prio-Only = 10%
Best Effort (be)	1	Rate = 100% CIR = 0% MBS = 5%	9	Rate = 100% CIR = 0% MBS = 5%

Forwarding class	Queue	Definition	Queue	Definition
		CBS = 0.1% High-Prio-Only = 10%		CBS = 0.1% High-Prio-Only = 10%

3.7.6 Network and service QoS queue parameters

The following queue parameters are provisioned on network and service queues:

- [Queue ID](#)
- [Committed information rate](#)
- [Peak information rate](#)
- [Adaptation rule](#)
- [Committed burst size](#)
- [Maximum burst size](#)
- [High-priority-only buffers](#)
- [High and low enqueueing thresholds](#)
- [Queue counters](#)
- [Queue type](#)
- [Queue mode](#)
- [Rate limiting](#)

3.7.6.1 Queue ID

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

3.7.6.2 Committed information rate

The CIR for a queue defines a limit for scheduling. Packets queued at service ingress queues are serviced by in-profile or out-of-profile schedulers based on the queue's CIR and the rate at which the packets are flowing. For each packet in a service ingress queue, the CIR is checked with the current transmission rate of the queue. If the current rate is at or below the CIR threshold, the transmitted packet is internally marked in-profile. If the flow rate is above the threshold, the transmitted packet is internally marked out-of-profile.

All 7705 SAR queues support the concept of in-profile and out-of-profile. The network QoS policy applied at network egress determines how or if the profile state is marked in packets transmitted into the network core. This is done by enabling or disabling the appropriate priority marking of network egress packets within a particular forwarding class. If the profile state is marked in the packets that are sent toward the network core, then out-of-profile packets are preferentially dropped over in-profile packets at congestion points in the network.

When defining the CIR for a queue, the value specified is the administrative CIR for the queue. The 7705 SAR maps a user-configured value to a hardware supported rate that it uses to determine the operational CIR for the queue. The user has control over how the administrative CIR is converted to an

operational CIR if a slight adjustment is required. The interpretation of the administrative CIR is discussed in [Adaptation rule](#).

The CIR value for a service queue is assigned to ingress and egress service queues based on service ingress QoS policies and service egress QoS policies, respectively.

The CIR value for a network queue is defined within a network queue policy specifically for the forwarding class. The *queue-id* parameter links the CIR values to the forwarding classes. The CIR values for the forwarding class queues are defined as a percentage of the network interface bandwidth.

3.7.6.3 Peak information rate

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

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

The PIR values for network queues are defined within network queue policies and are specific for each forwarding class. The PIR value for each queue for the forwarding class is defined as a percentage of the network interface bandwidth.

When defining the PIR for a queue, the value specified is the administrative PIR for the queue. The 7705 SAR maps a user-configured value to a hardware supported rate that it uses to determine the operational PIR for the queue. The user has control over how the administrative PIR is converted to an operational CIR if a slight adjustment is required. The interpretation of the administrative PIR is discussed in [Adaptation rule](#).

3.7.6.4 Adaptation rule

The schedulers on the network processor can only operate with a finite set of rates. These rates are called the operational rates. The configured rates for PIR and CIR do not necessarily correspond to the operational rates. In order to offer maximum flexibility to the user, the **adaptation-rule** command can be used to choose how an operational rate is selected based on the configured PIR or CIR rate.

The **max** parameter causes the network processor to be programmed at an operational rate that is less than the configured PIR or CIR rate by up to 1.0%. The **min** parameter causes the network processor to be programmed at an operational rate that is greater than the configured PIR or CIR rate by up to 1.0%. The **closest** parameter causes the network processor to be programmed at an operational rate that is closest to the configured PIR or CIR rate.

A 4-priority scheduler on the network processor of a third-generation (Gen-3) Ethernet adapter card or platform can be programmed at an operational CIR rate that exceeds 1.0% of the configured CIR rate. The PIR rate (that is, the maximum rate for the queue) and the SAP aggregate rates (CIR and PIR), maintain an accuracy of +/- 1.0% of the configured rates.

The average difference between the configured CIR rate and the programmed (operational) CIR rate is as follows:

- 2.0% for frame sizes that are less than 2049 bytes
- 4.0% for other frame sizes



Note: The percentages in the above list are averages only; the actual values can be higher.

The Gen-3 network processor PIR rate is programmed to an operational PIR rate that is within 1.0% of the configured rate, which ensures that the FC/CoS queue does not exceed its fair share of the total bandwidth.

3.7.6.5 Committed burst size

The CBS parameter specifies the committed buffer space allocated for a specific queue.

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

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

3.7.6.6 Maximum burst size

When the reserved buffers for a queue have been used, the queue contends with other queues for additional buffer resources up to the maximum burst size. The MBS parameter specifies the maximum queue depth to which a queue can grow. This parameter ensures that a traffic flow (that is, a customer or a traffic type within a customer port) that is massively or continuously oversubscribing the PIR of a queue will not consume all the available buffer resources. For high-priority forwarding class service queues, the MBS can be small because the high-priority service packets are scheduled with priority over other service forwarding classes. In other words, very small queues would be needed for high-priority traffic because the contents of the queues should have been scheduled by the best available scheduler.

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

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

3.7.6.7 High-priority-only buffers

High-priority-only buffers are defined on a queue and allow buffers to be reserved for traffic classified as high priority. When the queue depth reaches a specified level, only high-priority traffic can be enqueued. The high-priority-only reservation for a queue is defined as a percentage of the MBS value and has a default value of 10% of the MBS value.

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

On service egress, the high-priority-only reservation for a queue is defined in the service egress QoS policy. Service egress queues are specified by forwarding class. High-priority traffic for a given traffic class is traffic that has been marked as in-profile either on ingress classification or based on interpretation of the QoS markings.

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

3.7.6.8 High and low enqueueing thresholds

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

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

If an ingress queue is not configured with a **high-prio-only** parameter (the parameter is set to 0%), the low-priority and high-priority MBS thresholds are the same. There is no difference in high-priority and low-priority packet handling. At access ingress, the priority of a packet has no effect on which packets are scheduled first. Only the first buffering decision is affected. At ingress and egress, the current dynamic rate of the queue relative to the queue's CIR does affect the scheduling priority between queues going to the same destination (egress port).

From highest to lowest, the strict operating priority for queues is:

- expedited queues within the CIR (conform)
- best effort queues within the CIR (conform)
- expedited queues above the CIR (exceed)
- best effort queues above the CIR (exceed)

For access ingress, the CIR controls both dynamic scheduling priority and the marking threshold. At network ingress, the queue's CIR affects the scheduling priority but does not provide a profile marking function (as the network ingress policy trusts the received marking of the packet based on the network QoS policy).

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

3.7.6.9 Queue counters

The 7705 SAR maintains extensive counters for queues within the system to allow granular or extensive debugging and planning; that is, the usage of queues and the scheduler used for servicing a queue or packet is extremely useful in network planning activities. The following separate billing and accounting counters are maintained for each queue:

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

- counters for packets and octets transmitted out-of-profile

3.7.6.10 Queue type

The 7705 SAR allows two kinds of queue types: Expedited queues and Best Effort queues. Users can configure the queue type manually using the **expedite** and **best-effort** keywords, or automatically using the **auto-expedite** keyword. The queue type is specified as part of the **queue** command (for example, **config>qos>sap-ingress>queue queue-id queue-type create**).

With **expedite**, the queue is treated in an expedited manner, independent of the forwarding classes mapped to the queue.

With **best-effort**, the queue is treated in a non-expedited (best-effort) manner, independent of the forwarding classes mapped to the queue.

With **auto-expedite**, the queue type is automatically determined by the forwarding classes that are assigned to the queue. The queues that are set as **auto-expedite** are still either Expedited or Best Effort queues, but whether a queue is Expedited or Best Effort is determined by its assigned forwarding classes. In the default configuration, four of the eight forwarding classes (NC, H1, EF, and H2) result in an Expedited queue type, while the other four forwarding classes (L1, AF, L2, and BE) result in a Best Effort queue type.

Assigning one or more L1, AF, L2, and BE forwarding class to an Expedited queue results in a Best Effort queue type. See [Forwarding classes](#) for more information about default configuration values.

The **expedite**, **best-effort**, and **auto-expedite** queue types are mutually exclusive. Each defines the method that the system uses to service the queue from a hardware perspective.

3.7.6.11 Queue mode

The 7705 SAR supports two queue modes: priority mode and profile mode. Users can configure the queue mode using the **priority-mode** or **profile-mode** keywords when issuing the **config>qos>sap-ingress>queue queue-id queue-mode create** command. The default is **priority-mode**. The queue mode defines how an ingress access packet is categorized as in-profile or out-of-profile.

With priority mode, an access packet's in-profile or out-of-profile state is based on the dynamic rate of the ingress queue before being forwarded to the fabric. When the queue rate is lower than or equal to the configured CIR, the packet is considered in-profile. When the queue rate is higher than the CIR, the packet is considered out-of-profile. The profile state is determined when the packet is scheduled out of the queue, not when the packet is buffered into the queue.

With profile mode, the in-profile or out-of-profile state for packets assigned to a particular forwarding class is explicitly configured using the **config>qos>sap-ingress>fc>profile** command. This configuration places a forwarding class in a color-aware profile mode. Packets assigned to this forwarding class profile are only marked based on this profile marking if the forwarding class is mapped to a queue configured for **profile-mode**. If the forwarding class is mapped to a queue configured for **priority-mode**, the forwarding class profile setting is ignored and the packet's state is defined as in-profile or out-of-profile based on the dynamic rate of the ingress queue.

When the **profile in** command is executed on a forwarding class that is mapped to a queue operating in profile mode, all packets associated with the class are handled as in-profile. When the **profile out** command is executed on a forwarding class that is mapped to a queue operating in profile mode, all packets associated with the class are handled as out-of-profile.

When the **no profile** command is executed on a forwarding class that is mapped to a queue operating in profile mode, the data packets using the forwarding class are marked as in-profile or out-of-profile based on the dynamic rate of the ingress queue relative to its CIR.

Color-aware profiling adds the ability to selectively treat packets received on a SAP as in-profile (green) or out-of-profile (yellow) regardless of the queue forwarding rate. For example, a network operator can color a packet out-of-profile with the intention of preserving in-profile bandwidth for higher-priority packets.

A queue operating in profile mode can support in-profile, out-of-profile, and non-profiled packets simultaneously because multiple forwarding classes with different forwarding class profiles can be assigned to a single queue.

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

A SAP ingress queue operating in profile mode is classified as high-priority or low-priority based on the configuration of the forwarding class profile rather than on the high-priority or low-priority configuration specified for DSCP or dot1p. All non-profile packets flowing through the queue are considered high priority. Profiled in-profile packets are also handled as high priority, while profiled out-of-profile packets are handled as low priority.

For SAP ingress queues in profile mode, statistics are collected for color in (for a forwarding class configured for **profile in**), color out (for a forwarding class configured for **profile out**), and uncolor (for a forwarding class configured for **no profile**).

3.7.6.12 Rate limiting

The 7705 SAR supports egress-rate limiting and ingress-rate limiting on Ethernet ports.

The egress rate is set at the port level in the **config>port>ethernet** context.

Egress-rate limiting sets a limit on the amount of traffic that can leave the port to control the total bandwidth on the interface. If the egress-rate limit is reached, the port applies backpressure on the queues, which stops the flow of traffic until the queue buffers are emptied. This feature is useful in scenarios where there is a fixed amount of bandwidth; for example, a mobile operator who has leased a fixed amount of bandwidth from the service provider.

The **ingress-rate** command configures a policing action to rate-limit the ingress traffic. Ingress-rate enforcement uses dedicated hardware for rate limiting; however, software configuration is required at the port level (ingress-rate limiter) to ensure that the network processor or adapter card or port never receives more traffic than they are optimized for.

The configured ingress rate ensures that the network processor does not receive traffic greater than this configured value on a per-port basis. When the ingress-rate value is reached, all subsequent frames are dropped. The ingress-rate limiter drops excess traffic without determining whether the traffic has a higher or lower priority.

For more information about egress and ingress rate limiting, see the **egress-rate** and **ingress-rate** command descriptions in the 7705 SAR Interface Configuration Guide.

3.7.7 Slope policies (WRED and RED)

As part of 7705 SAR queue management, policies for WRED or RED queue management (also known as congestion management or buffer management) to manage the queue depths can be enabled at both access and network ports and associated with both ingress and egress queues. WRED policies can also be enabled on bridged domain (ring) ports.

Without WRED and RED, when a queue reaches its maximum fill size, the queue discards any new packets arriving at the queue (tail drop).

WRED and RED policies prevent a queue from reaching its maximum size by starting random discards when the queue reaches a user-configured threshold value. This avoids the impact of discarding all the new incoming packets. By starting random discards at this threshold, customer devices at an end-system may be adjusted to the available bandwidth.

As an example, TCP has built-in mechanisms to adjust for packet drops. TCP-based flows lower the transmission rate when some of the packets fail to reach the far end. This mode of operation provides a much better way of dealing with congestion than dropping all the packets after the whole queue space is depleted.

The WRED and RED curve algorithms are based on two user-configurable thresholds (minThreshold and maxThreshold) and a discard probability factor (maxDProbability) (see [Figure 24: WRED for high-priority and low-priority traffic in the same queue](#)). The minThreshold (minT) indicates the level when where discards start and the discard probability is zero. The maxThreshold (maxT) indicates the level where the discard probability reaches its maximum value. Beyond this the maxT level, all newly arriving packets are discarded. The steepness of the slope between minT and maxT is derived from the maxDProbability (maxDP). Therefore, the maxDP indicates the random discard probability at the maxT level.

The main difference between WRED and RED is that with WRED, there can be more than one curve managing the fill rate of the same queue.

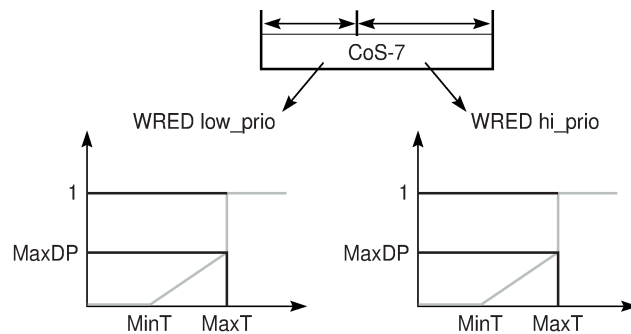
WRED slope curves can run against high-priority and low-priority traffic separately for ingress and egress queues. This allows the flexibility to treat low-priority and high-priority traffic differently. WRED slope policies are used to configure the minT, maxT and maxDP values, instead of configuring these thresholds against every queue. It is the slope policies that are then applied to individual queues. Therefore, WRED slope policies affect how and when the high-priority and low-priority traffic is discarded within the same queue.

Referring to the following figure, one WRED slope curve can manage discards on high-priority traffic and another WRED slope curve can manage discards on low-priority traffic. The minT, maxT and maxDP values configured for high-priority and low-priority traffic can be different and can start discarding traffic at different thresholds. The **start-avr**, **max-avr**, and **max-prob** commands are used to set the minThreshold, maxThreshold, and maxDProbability values, respectively.



Note: The figure shows a step function at maxT. The maxDP value is the target value entered for the configuration and it partly determines the slope of the weighting curve. At maxT, if the arrival of a new packet will overflow the buffer, the discard probability jumps to 1, which is not the maxDP value. Therefore, a step function exists in this graph.

Figure 24: WRED for high-priority and low-priority traffic in the same queue



19696

The formula to calculate the average queue size is:

$$\text{average queue size} = (\text{previous average} \times (1 - 1/2^{\text{TAF}})) + (\text{current queue size} \times 1/2^{\text{TAF}})$$

The Time Average Factor (TAF) is the exponential weight factor used in calculating the average queue size. The *time_average_factor* parameter is not user-configurable and is set to a system-wide default value of 3. By locking TAF to a static value of 3, the average queue size closely tracks the current queue size so that WRED can respond quickly to long queues.

3.7.7.1 WRED MinThreshold and MaxThreshold computation

CBS is configured in kilobytes through the CLI; MBS is configured in bytes or kilobytes. These configured values are converted to the corresponding number of buffers. The conversion factor is a non-user-configurable, fixed default value that is equal to the system-defined maximum frame size, ensuring that even the largest frames can be hosted in the allocated buffer pools. This type of WRED is called buffer-based WRED.

User-defined minThreshold and maxThreshold values, each defined as a percentage, are also converted to the number of buffers. The minT is converted to the system-minThreshold, and the maxT is converted to the system-maxThreshold.

The system-minT must be the absolute closest value to the minT that satisfies the formula below (2^x means 2 to the exponent x):

$$\text{system-maxThreshold} - \text{system-minThreshold} = 2^x$$



Note: The 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X Ethernet ports use payload-based WRED (also called byte-based WRED); see [Payload-based WRED](#). The above "system-minT" calculation does not apply to payload-based WRED. The 7705 SAR-X TDM ports use buffer-based WRED.

3.7.7.2 WRED on bridging domain (ring) queues

The bridging domain queues support the following DP (discard probability) values: 0% to 10%, 25%, 50%, 75%, and 100%. User-configured values are rounded down to match these DP values.

For example, configuring a DP to be 74% means that the actual value used is 50%.



Note: Tail drop for out-of-profile traffic begins when the queue occupancy of the out-of-profile traffic exceeds that of the committed buffer size.

3.7.7.3 Payload-based WRED

The third-generation Ethernet adapter cards and platforms use payload-based WRED instead of buffer-based WRED (see [WRED MinThreshold and MaxThreshold computation](#)). Payload-based WRED does not count the unused overhead space (empty space in the buffer) when making discard decisions, whereas buffer-based WRED counts the unused overhead space. Payload-based WRED is also referred to as byte-based WRED.

When a queue on an adapter card that uses payload-based WRED reaches its maximum fill (that is, the total byte count exceeds the configured maximum threshold), tail drop begins and operates in the same way as it does on any other adapter card or platform.

With payload-based WRED, the discard decision is based on the number of bytes in the queue instead of the number of buffers in the queue. For example, to accumulate 512 bytes of payload in a queue will take four buffers if the frame size is 128 bytes, but will take one buffer if the frame size is 512 bytes or more. Basing discards on bytes rather than buffers improves the efficient use of queues. In either case, byte- or buffer-based WRED, random discards begin at the minimum threshold (minT) point.

For example, assume a queue has MBS set to 512 kB (converts to 1000 buffers), minT (**start-avg**) is set to 10% (100 buffers), and maxT (**max-avg**) is set to 80% (800 buffers). The following table shows when discards and tail drop start when payload-based WRED is used.

Table 30: Payload-based WRED: discards and tail drop starts

Frame size	Discards start	Tail drop start
128 bytes	400 buffers in the queue (100 x 4)	3200 buffers in the queue (800 x 4)
512 bytes	100 buffers in the queue	800 buffers in the queue
1024 bytes	100 buffers in the queue	800 buffers in the queue

For tail drop, if the high-priority-only threshold is set to 10%:

- when any frame size is greater than or equal to 512 bytes, tail drop starts after 900 buffers are in use for low-priority traffic



Note:

- If an adapter card or platform other than a third-generation adapter card or platform is used in the previous example (that is, an adapter card or platform with buffer-based WRED), both WRED and tail drop start after 100 buffers are consumed, because both 128-byte and 512-byte frames fill one buffer (512 bytes).
- Because tail drop (which is buffer-based) and WRED (which is payload-based) operate differently, it is not recommended that tail drop and WRED be used in the same queue.

3.7.8 ATM traffic descriptor profiles

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

- sustained information rate (SIR)
- peak information rate (PIR)
- minimum information rate (MIR)
- maximum burst size (MBS)

QoS traffic descriptor profiles are applied on ATM VLL (Apipe) SAPs.

3.7.9 Fabric profiles

Fabric profiles allow access and network ingress to-fabric shapers to have user-configurable rates of switching throughput from an adapter card toward the fabric.

Two fabric profile modes are supported: per-destination mode and aggregate mode. Both modes offer shaping toward the fabric from an adapter card, but per-destination shapers offer the maximum flexibility by precisely controlling the amount of traffic to each destination card at a user-defined rate.

For the 7705 SAR-8 Shelf V2 and the 7705 SAR-18, the maximum rate depends on a number of factors, including platform, chassis variant, and slot type. See [Configurable ingress shaping to fabric \(access and network\)](#) for details. For information about fabric shaping on the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-Wx, see [Fabric shaping on the fixed platforms \(access and network\)](#).

3.7.10 Shaper policies

Shaper policies define dual-rate shaper parameters that control access or network traffic by providing tier-3 aggregate shaping to:

- shaped and unshaped SAP traffic for access ingress flows
- shaped and unshaped SAP traffic for access egress flows
- shaped and unshaped VLAN traffic for network egress flows



Note: For network egress traffic on a non-hybrid Gen-3 port, the CIR value of the shaper group is ignored because of the behavior of the 4-priority scheduler on Gen-3 hardware at network egress. For more information, see [QoS for Gen-3 adapter cards and platforms](#).

See [Per-SAP aggregate shapers \(H-QoS\) on Gen-2 hardware](#) and [Per-VLAN network egress shapers](#) for details on per-SAP and per-VLAN shapers.

3.7.11 QoS policy entities

Services are configured with default QoS policies. Additional policies must be explicitly created and associated. There is one default service ingress QoS policy, one default service egress QoS policy, and

one default network QoS policy. Only one ingress QoS policy and one egress QoS policy can be applied to a SAP or network port.

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

All QoS policy parameters can be configured in the CLI. QoS policies can be applied to the following service types:

- Epipe – both ingress and egress policies are supported on an Epipe SAP
- Apipe – both ingress and egress policies are supported on an Apipe SAP
- Cpipe – only ingress policies are supported on a Cpipe SAP
- Fpipe – both ingress and egress policies are supported on an Fpipe SAP
- Hpipe – both ingress and egress policies are supported on an Hpipe SAP
- Lpipe – both ingress and egress policies are supported on an Lpipe SAP

QoS policies can be applied to the following network entities:

- network ingress interface
- network egress interface

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

3.8 Configuration notes

The following guidelines and restrictions apply to the implementation of QoS policies:

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

4 Network QoS policies

This chapter provides information to configure network QoS policies using the command line interface (CLI).

Topics in this chapter include:

- [Overview](#)
- [Basic configuration](#)
- [Service management tasks](#)
- [Network QoS policy command reference](#)

4.1 Overview

The network QoS policy consists of an ingress and egress component for interfaces in the IP domain, and a ring component for interfaces in the bridging domain.

The ingress component of the QoS policy defines how DSCP bits (for GRE and IP) and multiprotocol label switching (MPLS) Experimental (EXP) bits are mapped to internal forwarding class and profile state. The forwarding class and profile state define the per-hop behavior (PHB) or the QoS treatment through the 7705 SAR.

The egress component of the QoS policy defines the DSCP bit, MPLS EXP bit, and the dot1p marking based on the forwarding class and the profile state.

The ring component of the QoS policy defines how dot1p bits are mapped to network queue and profile state.

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. Network policy-id 1 exists as the default policy that is applied to all network interfaces by default. The network policy-id 1 cannot be modified or deleted. For the ingress, it defines the default DSCP-to-FC and profile state, and MPLS EXP-to-FC and profile state mappings. For the egress, it defines eight forwarding classes and profile states that represent the packet marking criteria. For the ring, it defines eight dot1p-to-queue and profile state mappings.

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

4.2 Basic configuration

This section contains the following topics related to creating and applying network QoS policies for the IP domain and the bridging domain:

- [Configuring a network QoS policy](#)
- [Creating a network QoS policy](#)
- [Applying network QoS policies](#)
- [Default network QoS policy values](#)

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

- Each network QoS policy must have a unique policy ID.

4.2.1 Configuring a network QoS policy

Configuring and applying QoS policies other than the default policy is optional.

Define the following parameters to configure a network QoS policy (see [Default network QoS policy values](#) for default values):

- a network policy ID value – the system does not dynamically assign a value
- a network policy type – the type of network policy can be **ip-interface** for IP domain policies (to configure ingress and egress criteria), or **ring** for bridging domain policies (to configure ring criteria), or **default** (to set any network policy to its default values)
- a description – a text string description of policy features
- scope – the policy scope as exclusive or template
- egress criteria – you can modify egress criteria to customize the forwarding class to be instantiated. Otherwise, the default values are applied.
 - dot1p – the dot1p-in-profile and dot1p-out-profile mapping for the forwarding class
 - DSCP – the DSCP value is used for all GRE and IP packets (in or out of profile) requiring marking that egress on this forwarding class
 - LSP EXP – the EXP value is used for all MPLS labeled packets (in or out of profile) requiring marking that egress on this forwarding class
- ingress criteria – the DSCP to forwarding class mapping for all GRE and IP packets and the MPLS EXP bits to forwarding class mapping for all labeled packets
 - default-action – the default action to be taken for packets that have undefined DSCP or MPLS EXP bits set. The default-action specifies the forwarding class to which such packets are assigned. The default-action is automatically created when the policy is created.
 - DSCP – a mapping between the DSCP bits of the network ingress traffic and the forwarding class. Ingress traffic that matches the specified DSCP bits is assigned to the corresponding forwarding class.
 - LSP EXP – a mapping between the LSP EXP bits of the network ingress traffic and the forwarding class. Ingress traffic that matches the specified LSP EXP bits is assigned to the corresponding forwarding class.
- ring criteria – the dot1p-to-queue mapping for all packets arriving on a bridging domain port (that is, a ring port or the add/drop port)
 - default-action – the default action to be taken for packets that have undefined dot1p bits set. The default-action specifies the queue to which such packets are assigned and the profile state of the packets. The default-action is automatically created when the policy is created.
 - dot1p – a mapping of the dot1p bits of the ingress traffic to the queue and the profile state

4.2.2 Creating a network QoS policy

Configuring and applying network QoS policies other than the default policy is optional.

Use the following CLI syntax to create a network QoS policy for router interfaces (ip-interface type) and ring port (ring type).

The **ip-interface** keyword is optional for router interface network policies.

Use the **ring** keyword to create a network QoS policy that can be applied to a bridging domain port (that is, a ring port or the add/drop port). Up to eight dot1p-to-queue and profile mappings can be defined under the **ring** command.

CLI syntax:

```
config>qos#
  network network-policy-id [create] [network-policy-type {ip-interface
| ring | default}]
    description description-string
    scope {exclusive|template}
    egress
      fc {be|l2|af|l1|h2|ef|h1|nc}
        dot1p dot1p-priority
        dot1p-in-profile dot1p-priority
        dot1p-out-profile dot1p-priority
        dscp-in-profile dscp-name
        dscp-out-profile dscp-name
        lsp-exp-in-profile lsp-exp-value
        lsp-exp-out-profile lsp-exp-value
    ingress
      default-action fc {be|l2|af|l1|h2|ef|h1|nc} profile {in|out}
      dscp dscp-name fc {be|l2|af|l1|h2|ef|h1|nc} profile {in|out}
      ler-use-dscp
      lsp-exp lsp-exp-value fc fc-name profile {in|out}
    ring
      default-action queue queue-id profile {in | out}
      dot1p dot1p-priority queue queue-id profile {in|out}
  exit
```

Example:

```
configure qos network 700 create network-policy-type ip-interface
config>qos>network$ description "Net Policy"
config>qos>network$ scope template
config>qos>network$ egress fc be
config>qos>network>egress>fc$ dot1p 1
config>qos>network>egress>fc$ lsp-exp-in-profile 2
config>qos>network>egress>fc$ lsp-exp-out-profile 3
config>qos>network>egress>fc$ exit
config>qos>network$ ingress
config>qos>network>ingress$ default-action fc be profile in
config>qos>network>ingress$ exit
config>qos>network$ exit
```

The following output displays the configuration for an ip-interface type network policy 700:

```
*A:ALU-1>config>qos# info
-----
echo "QoS Policy Configuration"
#-----
network 700 network-policy-type ip-interface create
description "Net Policy"
```



```

    ingress
      default-action fc be profile in
    exit
    egress
      fc be
        lsp-exp-in-profile 2
        lsp-exp-out-profile 3
        dot1p-in-profile 1
        dot1p-out-profile 1
      exit
    exit
  exit
exit
-----

```

The following example creates a ring type network QoS policy on the 2-port 10GigE (Ethernet) Adapter card or 2-port 10GigE (Ethernet) module.

Example:

```

config>qos>network 5 network-policy-type ring create
config>qos>network>ring# default-action queue 1 profile out
config>qos>network>ring# dot1p 1 queue 4 profile in
config>qos>network>ring# dot1p 5 queue 6 profile out
config>qos>network>ring# exit
config>qos>network# scope template
config>qos>network# exit

```

The following output displays the configuration for ring network policy 5:

```

*A:7705custDoc:Sar18>config>qos>network# info detail
-----
    no description
    scope template
    ring
      default-action queue 1 profile out
      dot1p 1 queue 4 profile in
      dot1p 5 queue 6 profile out
    exit
  -----
*A:7705custDoc:Sar18>config>qos>network#

```

4.2.3 Applying network QoS policies

You can apply network QoS policies to router interfaces and a ring adapter card.

Use the following CLI syntax to apply network policies to router interfaces:

CLI syntax:

```

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

```

Example:

```

config>router# interface ALU-1
config>router>if$ qos 700
config>router>if$ exit

```

The following output displays the configuration for router interface ALU-1 with network policy 700 applied to the interface.

```
A:ALU-1>config>router# info
#-----
echo "IP Configuration"
#-----
      interface "ALU-1"
        qos 700
      exit
      interface "ip-10.0.0.2"
        address 10.10.0.2/16
      exit
#-----
```

Use the following CLI syntax to apply a ring type network policy to a ring adapter card. Applying the policy to the card means that the policy is applied to ingress traffic of the ring ports and the add/drop port:

CLI syntax:

```
config>card>mda>network
ring
qos-policy network-policy-id
```

Example:

```
config# card 1
config>card# mda 8
config>card>mda# network ring
config>card>mda>network>ring# qos-policy 5
```

The following output displays the configuration for ring type network policy 5 applied to a ring adapter card.

```
*A:7705custDoc:Sar18>config>card>mda>network# info detail
#-----
      ingress
        fabric-policy 1
        queue-policy "default"
      exit
      ring
        qos-policy 5
        add-drop-port-queue-policy "default"
      exit
#-----
*A:7705custDoc:Sar18>config>card>mda>network#
```

4.2.4 Default network QoS policy values

The default network policy is identified as policy-id 1. Default policies cannot be modified or deleted. [Table 25: Default network QoS policy DSCP-to-forwarding class mappings](#), [Table 26: Default network QoS policy LSP EXP-to-forwarding class mappings](#), and [Table 27: Default network QoS policy dot1p-to-queue class mappings](#) (found in the section on [Network and network queue QoS policies](#)) list the default network QoS policy parameters for ingress, egress, and ring policies.

The following output displays the default network policy configuration:

```
A:ALU-1>config>qos>network# info detail
#-----
```

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

```

        fc h2
            dscp-in-profile af41
            dscp-out-profile af42
            lsp-exp-in-profile 4
            lsp-exp-out-profile 4
            dot1p-in-profile 4
            dot1p-out-profile 4
        exit
        fc l1
            dscp-in-profile af21
            dscp-out-profile af22
            lsp-exp-in-profile 3
            lsp-exp-out-profile 2
            dot1p-in-profile 3
            dot1p-out-profile 3
        exit
    exit
ring
    default-action queue 1 profile out
    dot1p 0 queue 1 profile out
    dot1p 1 queue 2 profile in
    dot1p 2 queue 3 profile out
    dot1p 3 queue 3 profile in
    dot1p 4 queue 5 profile in
    dot1p 5 queue 6 profile in
    dot1p 6 queue 7 profile in
    dot1p 7 queue 8 profile in
exit
-----

```

4.3 Service management tasks

This section describes the following service management tasks:

- [Deleting QoS policies](#)
- [Copying and overwriting network policies](#)
- [Editing QoS policies](#)

4.3.1 Deleting QoS policies

A network policy is associated by default with router interfaces. You can replace the default policy with a non-default policy, but you cannot entirely remove the policy from the configuration. When you remove a non-default policy, the policy association reverts to the default network policy-id 1.

Use the following syntax to delete a network policy.

CLI syntax:

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

Example:

```
config>qos# no network700
```

4.3.2 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 an existing policy ID.

CLI syntax:

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

Example:

```
config>qos# copy network 1 600
config>qos# copy slope-policy 600 700
MINOR: CLI Destination "700" exists use {overwrite}.
config>qos# copy slope-policy 600 700 overwrite
config>qos#
```

The following output displays copied policies:

```
ALU-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
...
-----
ALU-12>config>qos#
```

4.3.3 Editing QoS policies

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

4.4 Network QoS policy command reference

4.4.1 Command hierarchies

- [Configuration commands](#)
 - [QoS policy network commands](#)
 - [Self-generated traffic configuration commands](#)
- [Operational commands](#)
- [Show commands](#)

4.4.1.1 Configuration commands

4.4.1.1.1 QoS policy network commands

```

config
- qos
- [no] network network-policy-id [create] [network-policy-type {ip-interface | ring
| default}]
- description description-string
- no description
- scope {exclusive | template}
- no scope
- egress
- [no] fc fc-name
- dot1p dot1p-priority
- no dot1p
- dot1p-in-profile dot1p-priority
- no dot1p-in-profile
- dot1p-out-profile dot1p-priority
- no dot1p-out-profile
- dscp-in-profile dscp-name
- no dscp-in-profile
- dscp-out-profile dscp-name
- no dscp-out-profile
- lsp-exp-in-profile lsp-exp-value
- no lsp-exp-in-profile
- lsp-exp-out-profile lsp-exp-value
- no lsp-exp-out-profile
- ingress
- default-action fc fc-name profile {in | out}
- dscp dscp-name fc fc-name profile {in | out}
- no dscp
- [no] ler-use-dscp
- lsp-exp lsp-exp-value fc fc-name profile {in | out}
- no lsp-exp lsp-exp-value
- ring
- default-action queue queue-id profile {in | out}
- dot1p dot1p-priority queue queue-id profile {in | out}
- no dot1p dot1p-priority

```

4.4.1.1.2 Self-generated traffic configuration commands

```
config
- router
- sgt-qos
- application dscp-app-name dscp {none | dscp-value | dscp-name} [fc-queue fc-name]
profile {in | out}}
- application dot1p-app-name dot1p {none | dot1p-priority} [fc-queue fc-name]
profile {in | out}}
- no application {dscp-app-name | dot1p-app-name}
- dscp dscp-name fc fc-name
- no dscp dscp-name
```

4.4.1.2 Operational commands

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

4.4.1.3 Show commands

```
show
- qos
- dscp-table [value dscp-value]
- network policy-id [detail]
show
- router
- sgt-qos
- application [app-name] [dscp | dot1p]
- dscp-map [dscp-name]
```

4.4.2 Command descriptions

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

4.4.2.1 Configuration commands

- [Generic commands](#)
- [Network QoS policy commands](#)
- [Network egress QoS policy commands](#)
- [Network ingress QoS policy commands](#)
- [Network ring QoS policy commands](#)
- [Network egress QoS policy forwarding class commands](#)
- [Self-generated traffic commands](#)

4.4.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>network

config>qos>mc-mlppp>sap-egress

Description

This command associates a text string with a configuration context to help identify the context in the configuration file.

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

Default

n/a

Parameters

description-string

a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

4.4.2.1.2 Network QoS policy commands

network

Syntax

[no] **network** *network-policy-id* [create] [**network-policy-type** {**ip-interface** | **ring** | **default**}]

Context

config>qos

Description

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

The **network-policy-type** keyword defines the type of network policy that will be created. The **ip-interface** type network policy is assigned to router interfaces. The **ring** type network policy is assigned to bridging domain ports on a ring adapter card.

Using the **network-policy-type** keyword is optional. If **network-policy-type** is not used, a default ip-interface policy is created.



Note: With the addition of ring ports on the 2-port 10GigE (Ethernet) Adapter card, the **network-policy-type** keyword allows the creation of network QoS policies for IP interfaces in the IP domain (that is, ingress and egress network ports) and for ring policies in the bridging domain.

Network *policy-id* 1 exists as the default policy that is applied to all network interfaces by default. Network *policy-id* 1 cannot be modified or deleted, and is reapplied when the **network-policy-type default** keyword is used.

If a new network policy is created for the IP domain (for instance, *policy-id* 2), only the default action and egress forwarding class parameters are identical to the default *policy-id* 1. A new network policy does not contain the default DSCP-to-FC or MPLS EXP-to-FC mapping. To create a new network policy that includes the default ingress DSCP-to-FC or MPLS EXP-to-FC mapping, the default network *policy-id* 1 can be copied (using the **copy** command). You can modify parameters or use the **no** modifier to remove an object from the configuration. Similarly, the copy and modify process is used to create a new ring type network policy, where the dot1p-to-queue and profile state mapping are available for use in the new policy.

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

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

Default

System Default Network Policy 1 defined

Parameters

network-policy-id

uniquely identifies the policy on the 7705 SAR

Values 1 to 65535

Default 1

create

keyword used to create a network QoS policy

network-policy-type {ip-interface | ring | default}

keyword used to define the type of network policy

Values **ip-interface**: creates a network QoS policy for the IP domain by providing access to the [egress](#) and [ingress](#) commands
ring: creates a ring network QoS policy for the bridging domain by providing access to the [ring](#) command
default: sets the network policy type to policy ID 1

scope

Syntax

scope {exclusive | template}

no scope

Context

config>qos>network

Description

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

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

Default

template

Parameters

exclusive

when the scope of a policy is defined as exclusive, the policy can only be applied to one network. If a policy with an exclusive scope is assigned to a second network, an error message is generated. If the policy is removed from the exclusive network, it will become available for assignment to another exclusive network.

The system default policies cannot be defined as exclusive scope. An error will be generated if scope exclusive is executed in any policies with a *policy-id* equal to 1.

template

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

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

4.4.2.1.3 Network egress QoS policy commands

egress

Syntax

egress

Context

config>qos>network

Description

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

The forwarding class and profile state mapping to in-profile and out-of-profile DSCP and MPLS EXP bits mapping for all labeled packets are also defined under this node.

For MPLS tunnels, if network egress Ethernet ports are used, dot1p bit marking can be enabled in conjunction with EXP bit marking. In this case, the tunnel and pseudowire EXP bits do not have to be the same as the dot1p bits.

For GRE and IP tunnels, dot1p marking and pseudowire EXP marking can be enabled, and DSCP marking can also be enabled.

The service packets are transported over an MPLS LSP, GRE tunnel, or IP tunnel.

All out-of-profile service packets are marked with the corresponding DSCP (for GRE or IP packets) or EXP (for MPLS packets) bit value at network egress. All in-profile service packets are marked with the corresponding in-profile DSCP or EXP bit value based on the forwarding class they belong to.

fc

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

Context
config>qos>network>egress

Description
This command specifies the forwarding class name. The **fc** *fc-name* represents a CLI parent node that contains sub-commands or parameters describing the marking criteria for that forwarding class. The **fc** command overrides the default parameters for that forwarding class defined in the network default *policy-id* 1.
The **no** form of this command reverts to the defined parameters in the default network policy *policy-id* 1. If the *fc-name* is removed from the default network policy *policy-id* 1, that forwarding class reverts to the factory defaults.

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

Parameters
fc-name
the case-sensitive, system-defined forwarding class name for which policy entries will be created

Values	be, l2, af, l1, h2, ef, h1, nc
Default	n/a

4.4.2.1.4 Network ingress QoS policy commands

ingress

Syntax
ingress

Context
config>qos>network

Description

This command is used to enter the CLI mode that creates or edits policy entries that specify the DSCP to forwarding class mapping for all GRE or IP packets and define the MPLS EXP bits to forwarding class mapping for all labeled packets.

When pre-marked GRE, IP, or MPLS packets ingress on a network port, they get a per-hop behavior (that is, the QoS treatment through the 7705 SAR based on the mapping defined under the current node).

default-action

Syntax

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

Context

config>qos>network>ingress

Description

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

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

Default

n/a

Parameters

fc-name

specifies the forwarding class name. All packets with DSCP or MPLS EXP bits not defined will be placed in this forwarding class.

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

Default be

profile {in | out}

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

Values in, out

Default out

dscp

Syntax

dscp *dscp-name* **fc** *fc-name* **profile** {**in** | **out**}

no dscp

Context

config>qos>network>ingress

Description

This command creates a mapping between the DSCP of the network ingress traffic and the forwarding class for GRE or IP packets.

Ingress traffic that matches the specified DSCP is assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all 64 DSCP values with the forwarding class. For undefined code points, packets are assigned to the forwarding class specified under the default-action command.

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

Default

n/a

Parameters

dscp-name

specifies the DSCP to be associated with the forwarding class. The DSCP value is derived from the most significant six bits in the IP header ToS byte field (DSCP bits). The six DSCP bits define 64 DSCP values used to map packets to per-hop QoS behavior.

A maximum of 64 DSCP rules are allowed on a single policy. The specified name must exist as a *dscp-name*. [Table 31: Valid DSCP names](#) lists all the valid DSCP names.

Table 31: Valid DSCP names

dscp-name
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

fc-name

specifies the forwarding class name with which the DSCP will be associated

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

Default be

profile {in | out}

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

Values in, out

Default out

ler-use-dscp

Syntax

[no] ler-use-dscp

Context

config>qos>network>ingress

Description

This command is used to override tunnel QoS mapping on all ingress network IP interfaces that the *network-policy-id* is associated with. The command may be defined at any time after the network QoS policy has been created.

For IP traffic riding over MPLS or GRE tunnels that will be routed to the base router, a VPRN interface, or an IES interface at the tunnel termination point (the eLER), this command makes it possible for the 7705 SAR to ignore the EXP/DSCP bits in the tunnel header when the packets arrive at the eLER. This is useful when the mapping for the tunnel QoS marking does not completely reflect the required QoS handling for the IP routed packet. When the command is enabled on an ingress network IP interface, the IP interface will ignore the tunnel's QoS mapping and will derive the internal forwarding class and associated profile state based on the DSCP values of the IP header ToS field. This command applies only on the eLER where the tunnel or service is terminated and the next header in the packet is IP.

The default state is to not enforce customer DSCP at the tunnel-termination point, IP-routed QoS override within the network QoS policy.

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

Default

no ler-use-dscp

lsp-exp

Syntax

lsp-exp *lsp-exp-value* **fc** *fc-name* **profile** {in | out}

no lsp-exp *lsp-exp-value*

Context

config>qos>network>ingress

Description

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

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

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

Default

n/a

Parameters

lsp-exp-value

specifies the LSP EXP values to be associated with the forwarding class

Values 0 to 7 (decimal representation of 3-bit EXP field)

Default n/a

fc-name

specifies the FC name that the EXP bit pattern will be associated with

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

Default n/a

profile {in | out}

indicates whether the LSP EXP value is the in-profile or out-of-profile value

Values in, out

Default out

4.4.2.1.5 Network ring QoS policy commands

ring

Syntax

ring

Context

config>qos>network

Description

This command is used to enter the CLI context that creates or edits policy entries that specify the dot1p-to-queue mapping for all packets.

default-action

Syntax

default-action queue *queue-id* profile {in | out}

Context

config>qos>network>ring

Description

This command defines or edits the default action to be taken for packets that have undefined dot1p bits set. The **default-action** command specifies the queue to which received packets are assigned as well as their profile state. Multiple **default-action** commands will overwrite each previous **default-action** command.

Default

n/a

Parameters

queue-id

specifies the queue ID. All packets with dot1p bits not defined by a **dot1p** command will be placed in this queue.

Values 1 to 8

Default 1

profile {in | out}

all packets with dot1p bits not defined by a **dot1p** command will be considered in-profile or out-of-profile based on this command. In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets.

Values in, out

Default out

dot1p

Syntax

```
dot1p dot1p-priority queue queue-id profile {in | out}  
no dot1p dot1p-priority
```

Context

```
config>qos>network>ring
```

Description

This command creates a mapping between the dot1p bits of the ingress traffic and the queue and profile state. A maximum of eight dot1p entries are allowed on a single policy.

The **no** form of the command removes the association of the dot1p bit value with the queue and profile state. The default-action then applies to that dot1p bit pattern.

Default

```
0
```

Parameters

dot1p-priority
specifies the dot1p bit value to be associated with the queue and the profile state

Values	0 to 7
Default	n/a

queue-id
specifies the queue in which all packets with the dot1p bits value are placed

Values	1 to 8
Default	1

profile {in | out}
specifies the profile state of packets with the dot1p bit value, either in-profile or out-of-profile

Values	in, out
Default	out


4.4.2.1.6 Network egress QoS policy forwarding class commands

dot1p

Syntax
`dot1p dot1p-priority`
`no dot1p`

Context
config>qos>network>egress>fc

Description
This command explicitly defines the egress dot1p priority bits values for the forwarding class.



Note: When a single *dot1p-priority* is specified, it is applied to both in-profile and out-of-profile packets. The other forms of the command described below (**dot1p-in-profile** and **dot1p-out-profile**) allow different dot1p values for in-profile or out-of-profile packets to be specified.

The **no** form of the command sets the dot1p priority bits value to 0.

Default
0

Parameters
dot1p-priority
the explicit dot1p value for the specified forwarding class. Setting the value to 0 is equivalent to removing the marking value.

Values 0 to 7

Default n/a

dot1p-in-profile

Syntax
`dot1p-in-profile dot1p-priority`
`no dot1p-in-profile`

Context
config>qos>network>egress>fc

Description

This command specifies dot1p in-profile mappings.

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

Parameters

dot1p-priority

defines the dot1p marking for the forwarding class

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

Description

This command specifies dot1p out-profile mappings.

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

Parameters

dot1p-priority

defines the dot1p marking for the forwarding class

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

Description

This command specifies the in-profile DSCP name for the forwarding class. The corresponding DSCP value is used for all in-profile GRE or IP packets that require marking at egress on this forwarding class.

IP traffic from network interfaces is always trusted. There is no re-marking performed on network egress for global routing table (GRT) forwarded IP traffic (with the exception of GRE and IPSec tunnels).

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

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

Default

policy-id 1: factory setting

policy-id 2 to 65535: *policy-id* 1 setting

Parameters

dscp-name

specifies the DSCP to be associated with the forwarding class. The DSCP value is derived from the most significant six bits in the IP header ToS byte field (DSCP bits). The six DSCP bits define 64 DSCP values used to map packets to per-hop QoS behavior.

A maximum of 64 DSCP rules are allowed on a single policy. The specified name must exist as a *dscp-name*. [Table 31: Valid DSCP names](#) lists all the valid DSCP names.

dscp-out-profile

Syntax

dscp-out-profile *dscp-name*

no dscp-out-profile

Context

config>qos>network>egress>fc

Description

This command specifies the out-of-profile DSCP name for the forwarding class. The corresponding DSCP value is for all out-of-profile GRE or IP packets that require marking at egress on this forwarding class.

IP traffic from network interfaces is always trusted. There is no re-marking performed on network egress for global routing table (GRT) forwarded IP traffic (with the exception of GRE and IPSec tunnels).

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

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

Default

policy-id 1: factory setting

policy-id 2 to 65535: *policy-id* 1 setting

Parameters

dscp-name

specifies the DSCP to be associated with the forwarding class. The DSCP value is derived from the most significant six bits in the IP header ToS byte field (DSCP bits). The six DSCP bits define 64 DSCP values used to map packets to per-hop QoS behavior.

A maximum of 64 DSCP rules are allowed on a single policy. The specified name must exist as a *dscp-name*. [Table 31: Valid DSCP names](#) lists all the valid DSCP names.

lsp-exp-in-profile

Syntax

lsp-exp-in-profile *lsp-exp-value*

no lsp-exp-in-profile

Context

config>qos>network>egress>fc

Description

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

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

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

Default

policy-id 1: factory setting

policy-id 2 to 65535: *policy-id* 1 setting

Parameters

lsp-exp-value

the 3-bit LSP EXP bit value, expressed as a decimal integer

Values	0 to 7
Default	n/a

lsp-exp-out-profile

Syntax

lsp-exp-out-profile *lsp-exp-value*

no lsp-exp-out-profile

Context

```
config>qos>network>egress>fc
```

Description

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

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

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

Default

policy-id 1: factory setting

policy-id 2 to 65535: *policy-id* 1 setting

Parameters

lsp-exp-value

the 3-bit LSP EXP bit value, expressed as a decimal integer

Values 0 to 7

Default n/a

4.4.2.1.7 Self-generated traffic commands**sgt-qos****Syntax**

sgt-qos

Context

```
config>router
```

Description

This command enables the context to configure DSCP or dot1p re-marking for self-generated traffic (SGT).

application**Syntax**

application *dscp-app-name* **dscp** {**none** | *dscp-value* | *dscp-name*} [**fc-queue** *fc-name* **profile** {**in** | **out**}]

application *dot1p-app-name* **dot1p** {**none** | *dot 1p-priority*} [**fc-queue** *fc-name* **profile** {**in** | **out**}]

no application {*dscp-app-name* | *dot1p-app-name*}

Context

```
config>router>sgt-qos
```

Description

This set of commands configures DSCP marking for self-generated IP traffic or dot1p marking for self-generated non-IP traffic (specifically, IS-IS and ARP traffic).

When an IP or Layer 3 application is configured using the *dscp-app-name* parameter, the specified DSCP name or DSCP value is used for all packets generated by this application within the router instance in which it is configured. The value set in this command sets the DSCP value in the egress IP header. The egress QoS policy will not overwrite this value.

When a Layer 2 application is configured using the *dot1p-app-name* parameter, the specified dot1p priority value is used for all packets generated by this application within the router instance in which it is configured.

Only one name or value can be configured per application. If multiple entries are configured, a subsequent entry overrides the previously configured entry.

The **fc-queue** option redirects SGT applications to egress data queues rather than the default control queue by assigning them to a forwarding class. If this option is configured, the profile state must be set. All packets that are assigned to this forwarding class will be considered in-profile or out-of-profile based on the configuration. In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets.

If the **fc-queue** option is used with the *dscp-app-name* application, any configuration done using the **sgt-qos>dscp** command is ignored for packets generated by this application, as illustrated in the following examples:

```
sgt-qos>application telnet dscp cp1
```

```
sgt-qos>dscp cp1 fc af
```

```
sgt-qos>application ftp dscp cp1 fc-queue be profile out
```

```
sgt-qos>dscp cp1 fc af
```

In the first example, all packets generated by the Telnet application use DSCP CP1 and map to FC AF as configured with the **dscp** command. The dot1p bits of the outgoing packets are marked from the value that FC AF points to in the egress QoS policy.

In the second example, all packets generated by the FTP application use DSCP CP1 and map to FC BE as dictated by the **fc-queue** redirection. The dot1p bits of the outgoing packets are marked from the value that FC BE points to in the egress QoS policy. Because redirection is configured, the mapping configured with the **dscp** command is ignored.



Note: The above behavior applies to all SGT IP applications with the exception of VRRP, where the dot1p value is always set to 7, regardless of the value in the FC egress QoS policy.

If the **fc-queue** option is used with the *dot1p-app-name* application, the dot1p bits of the outgoing packets are marked with the value set with the *dot1p-priority* parameter, regardless of the value in the FC egress queue policy.

The **no** form of this command resets the DSCP or dot1p value for the application to its default value and resets the application to use the egress control queue.

Default

n/a

Parameters

dscp-app-name

the DSCP application name

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



Note:

- PTP in the context of SGT QoS is defined as Precision Timing Protocol and is an application in the 7705 SAR. The PTP application name is also used in areas such as event-control and logging. Precision Timing Protocol is defined in IEEE 1588-2008.
- PTP in the context of IP filters is defined as Performance Transparency Protocol. IP protocols can be used as IP filter match criteria; the match is made on the 8-bit protocol field in the IP header.

dscp-value

the value that maps to the DSCP name (the value **none** specifies that the default DSCP value for the application be used; see [Table 28: Applications and support for configurable DSCP or dot1p markings](#))

Values none | 0 to 63

dscp-name

the DSCP to be associated with the forwarding class. [Table 31: Valid DSCP names](#) lists the valid DSCP names.

dot1p-app-name

the dot1p application name

Values arp, isis

dot1p-priority

the dot1p priority (the value **none** specifies that the default dot1p value for the application be used; see [Table 28: Applications and support for configurable DSCP or dot1p markings](#))

Values none | 0 to 7

fc-name

the forwarding class assigned to SGT applications redirected to data queues

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


profile {in | out}
the profile state of packets assigned to the specified forwarding class; this parameter must be specified when the **fc-queue** parameter is configured

dscp

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

Context
config>router>sgt-qos


Description
This command creates a mapping between the DSCP of the self-generated traffic and the forwarding class. The forwarding class dot1p network QoS policy mapping is used to mark the dot1p bits of the Layer 3 or IP application. For example, configuring the *dscp-name* parameter as **be** and the *fc-name* parameter as **I1** results in marking the dot1p bits of the outgoing Ethernet frame, which is transporting self-generated IP traffic with DSCP bits set to BE, to the value that FC L1 points to in the network QoS policy (as configured in the **config>qos>network>egress>fc** context).

 **Note:** The dot1p class of service may not apply to all IP traffic and is dependent on the egress port encapsulation type.

Based on this configured FC, the network QoS policy for the egress forwarding complex sets the IEEE 802.1 dot1p bits.

Multiple commands can be entered to associate some or all of the 64 DSCP values with the forwarding class. For undefined code points, packets are assigned to the default forwarding class for the DSCP value. This value can be seen in the **show router sgt-qos dscp-map** output under the Default FC Value column.

The **no** form of the command resets the DSCP value to its default forwarding class.

 **Note:** If the **fc-queue** option is configured in the **sgt-qos>application dscp-app-name** command, the mapping created with this command is ignored for packets generated by the applications that are configured with the option.

The following table lists the default FC value for each DSCP value.

Table 32: DSCP-to-default FC value mapping

DSCP value	Default FC value
be	nc
cp1	be
cp2	be
cp3	be

DSCP value	Default FC value
cp4	be
cp5	be
cp6	be
cp7	be
cs1	be
cp9	be
af11	af
cp11	be
af12	af
cp13	be
af13	af
cp15	be
cs2	be
cp17	be
af21	l1
cp19	be
af22	l1
cp21	be
af23	l1
cp23	be
cs3	be
cp25	be
af31	l1
cp27	be
af32	l1
cp29	be
af33	l1

DSCP value	Default FC value
cp31	be
cs4	be
cp33	be
af41	nc
cp35	be
af42	h2
cp37	be
af43	h2
cp39	be
cs5	be
cp41	be
cp42	be
cp43	be
cp44	be
cp45	be
ef	ef
cp47	be
nc1	nc
cp49	be
cp50	h2
cp51	be
cp52	be
cp53	be
cp54	be
cp55	be
nc2	nc
cp57	be

DSCP value	Default FC value
cp58	be
cp59	be
cp60	be
cp61	be
cp62	be
cp63	be

Default
Listed in table

Parameters

dscp-name
the DSCP to be associated with the forwarding class. [Table 31: Valid DSCP names](#) lists the valid DSCP names.

fc-name
the forwarding class name with which the DSCP will be associated

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

4.4.2.2 Operational commands

copy

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

Context
config>qos

Description
This command copies existing QoS policy entries for a QoS policy ID to another QoS policy ID.
The **copy** command is used to create new policies using existing policies and also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters

src-pol dst-pol

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

Values 1 to 65535

overwrite


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

NOK>config>qos# copy network 1 427

MINOR: CLI Destination "427" exists use {overwrite}.

NOK>config>qos# copy network 1 427 overwrite

4.4.2.3 Show commands



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

dscp-table

Syntax

dscp-table [value *dscp-value*]

Context

show>qos

Description

This command displays DSCP name to DSCP value mappings.

Parameters

value *dscp-value*

the specific DSCP value for which to display information

Values 0 to 63

Default show all values

Output

The following output is an example of DSCP name to DSCP value mappings information, and [Table 33: DSCP name-to-value mapping field descriptions](#) describes the fields.

Output example

```
*A:ALU-1# show qos dscp-table
```

```
=====
```

DSCP Mapping

```
=====
```

DSCP Name	DSCP Value	TOS (bin)	TOS (hex)
be	0	0000 0000	00
cp1	1	0000 0100	04
cp2	2	0000 1000	08
cp3	3	0000 1100	0C
cp4	4	0001 0000	10
cp5	5	0001 0100	14
cp6	6	0001 1000	18
cp7	7	0001 1100	1C
cs1	8	0010 0000	20
cp9	9	0010 0100	24
af11	10	0010 1000	28
cp11	11	0010 1100	2C
af12	12	0011 0000	30
cp13	13	0011 0100	34
af13	14	0011 1000	38
cp15	15	0011 1100	3C
cs2	16	0100 0000	40
cp17	17	0100 0100	44
af21	18	0100 1000	48
cp19	19	0100 1100	4C
af22	20	0101 0000	50
cp21	21	0101 0100	54
af23	22	0101 1000	58
cp23	23	0101 1100	5C
cs3	24	0110 0000	60
cp25	25	0110 0100	64
af31	26	0110 1000	68
cp27	27	0110 1100	6C
af32	28	0111 0000	70
cp29	29	0111 0100	74
af33	30	0111 1000	78
cp31	31	0111 1100	7C
cs4	32	1000 0000	80
cp33	33	1000 0100	84
af41	34	1000 1000	88
cp35	35	1000 1100	8C
af42	36	1001 0000	90
cp37	37	1001 0100	94
af43	38	1001 1000	98
cp39	39	1001 1100	9C
cs5	40	1010 0000	A0
cp41	41	1010 0100	A4
cp42	42	1010 1000	A8
cp43	43	1010 1100	AC
cp44	44	1011 0000	B0
cp45	45	1011 0100	B4
ef	46	1011 1000	B8
cp47	47	1011 1100	BC
nc1	48	1100 0000	C0
cp49	49	1100 0100	C4
cp50	50	1100 1000	C8
cp51	51	1100 1100	CC
cp52	52	1101 0000	D0
cp53	53	1101 0100	D4
cp54	54	1101 1000	D8
cp55	55	1101 1100	DC

nc2	56	1110 0000	E0
cp57	57	1110 0100	E4
cp58	58	1110 1000	E8
cp59	59	1110 1100	EC
cp60	60	1111 0000	F0
cp61	61	1111 0100	F4
cp62	62	1111 1000	F8
cp63	63	1111 1100	FC
=====			
*A:ALU-1#			

Table 33: DSCP name-to-value mapping field descriptions

Label	Description
DSCP Name	The name of the DSCP to be associated with the forwarding class
DSCP Value	The DSCP value ranges (from 0 to 63)
TOS (bin)	The type of service in binary format
TOS (hex)	The type of service in hexadecimal format

network

Syntax

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

Context

show>qos

Description

This command displays network policy information.

Parameters

policy-id

displays information for the specific policy ID

Values 1 to 65535

Default all network policies

detail

displays detailed information for the specific policy ID

Output

The following outputs are examples of network policy information:

- Network policy information ([Output example](#), [Table 34: Network policy field descriptions](#))
- Ethernet ring policy information ([Output example](#), [Table 35: Ethernet ring network policy field descriptions](#))

Output example

```
*A:ALU-1# show qos network
=====
Network Policies
=====
Policy-Id      Description
-----
1              Default network QoS policy.
100           Network QoS policy 100.
=====

*A:ALU-1# show qos network detail
=====
QoS Network Policy
=====
Network Policy (1)
-----
Policy-id      : 1
Forward Class  : be                      Profile      : Out
Ring Queue     : 1                      Ring Profile  : Out
Scope         : Template                Policy Type   : Default
Description    : Default network QoS policy.
-----
DSCP           Forwarding Class      Profile
-----
be             be              Out
ef             ef              In
cs1            l2              In
nc1            h1              In
nc2            nc              In
af11           af              In
af12           af              Out
af13           af              Out
af21           l1              In
af22           l1              Out
af23           l1              Out
af31           l1              In
af32           l1              Out
af33           l1              Out
af41           h2              In
af42           h2              Out
af43           h2              Out
-----
LSP EXP Bit Map      Forwarding Class      Profile
-----
0                    be              Out
1                    l2              In
2                    af              Out
3                    af              In
4                    h2              In
5                    ef              In
6                    h1              In
7                    nc              In
-----
Dot1p Bit Map        Queue          Profile
-----
```

0	1	Out
1	2	In
2	3	Out
3	3	In
4	5	In
5	6	In
6	7	In
7	8	In

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	
- LSP EXP Bit 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	
- LSP EXP Bit 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 : 2	
- LSP EXP Bit 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 : 3	In-Profile : 3	
- LSP EXP Bit Mapping		
Out-of-Profile : 2	In-Profile : 3	
FC Value : 4	FC Name : h2	
- DSCP Mapping		
Out-of-Profile : af42	In-Profile : af41	
- Dot1p Mapping		
Out-of-Profile : 4	In-Profile : 4	
- LSP EXP Bit Mapping		
Out-of-Profile : 4	In-Profile : 4	
FC Value : 5	FC Name : ef	
- DSCP Mapping		
Out-of-Profile : ef	In-Profile : ef	
- Dot1p Mapping		
Out-of-Profile : 5	In-Profile : 5	
- LSP EXP Bit Mapping		
Out-of-Profile : 5	In-Profile : 5	
FC Value : 6	FC Name : h1	
- DSCP Mapping		
Out-of-Profile : nc1	In-Profile : nc1	
- Dot1p Mapping		
Out-of-Profile : 6	In-Profile : 6	
- LSP EXP Bit Mapping		
Out-of-Profile : 6	In-Profile : 6	
FC Value : 7	FC Name : nc	
- DSCP Mapping		
Out-of-Profile : nc2	In-Profile : nc2	
- Dot1p Mapping		

```

Out-of-Profile : 7
- LSP EXP Bit Mapping
Out-of-Profile : 7
In-Profile : 7
In-Profile : 7
-----
Interface Association
-----
Interface      : system
IP Addr.       : n/a
Interface      : address
IP Addr.       : n/a
Interface      : back
IP Addr.       : n/a
Interface      : dhcp_interface
IP Addr.       : n/a
Interface      : int_formp1s
IP Addr.       : n/a
Interface      : interface_1
IP Addr.       : n/a
Interface      : router_interface_1
IP Addr.       : 192.168.0.0/16
Interface      : vprn_interface
IP Addr.       : n/a
Interface      : ipv6_interface
IP Addr.       : n/a
Interface      : management
IP Addr.       : 192.168.0.10/16
Port Id       : system
Port Id       : n/a
Port Id       : n/a
Port Id       : n/a
Port Id       : n/a
Port Id       : 1/3/11.1
Port Id       : n/a
Port Id       : n/a
Port Id       : n/a
Port Id       : n/a
Port Id       : A/1
-----
Ring MDA Association
-----
MDA            : 1/11
=====

```

Table 34: Network policy field descriptions

Label	Description
Policy-Id	The ID that uniquely identifies the policy
Forward Class	The forwarding class name
Profile	The profile state of the traffic: In or Out
Ring Queue	The queue identifier
Ring Profile	The profile state of the ring traffic: In or Out
Scope	The scope of the policy: Template or Exclusive
Policy Type	The type of network policy: Ip-interface, Ring, or Default
Description	A text string that helps identify the policy's context in the configuration file
DSCP	The DSCP name associated with the forwarding class: Forwarding Class – the forwarding class associated with the DSCP Profile – whether the DSCP mapping pertains to in-profile or out-of-profile traffic

Label	Description
LSP EXP Bit Map	<p>The LSP EXP mapping value used for in-profile or out-of-profile traffic:</p> <p>Forwarding Class – the default-action forwarding class name. All packets with MPLS EXP bits not defined will be placed in this forwarding class.</p> <p>Profile – whether the LSP EXP bit mapping pertains to in-profile or out-of-profile traffic</p>
Dot1p Bit Map	<p>The dot1p mapping value used for queue and profile state:</p> <p>Queue – the queue in which all packets with the associated dot1p bit values will be placed</p> <p>Profile – whether the dot1p bit mapping pertains to in-profile or out-of-profile traffic</p>
Egress/Ingress Forwarding Class Queuing	
FC Value	The forwarding class value
FC Name	The forwarding class name
DSCP Mapping	Out-of-Profile - the out-of-profile DSCP mapping for the forwarding class
	In-Profile - the in-profile DSCP mapping for the forwarding class
Dot1p Mapping	Out-of-Profile - the out-of-profile dot1p bit mapping for the forwarding class
	In-Profile - the in-profile dot1p bit mapping for the forwarding class
LSP EXP Bit Mapping	Out-of-Profile - the out-of-profile LSP EXP bit mapping for the forwarding class
	In-Profile - the in-profile LSP EXP bit mapping for the forwarding class
Interface Association	
Interface	The name of the interface
IP Addr.	The IP address of the interface
Port Id	The physical port identifier that associates the interface
Ring MDA Association	
MDA	The CLI identifier of the adapter card associated with the policy

Output example

```

show qos network <policy-id>
=====
QoS Network Policy
=====
-----
Network Policy (1)
Policy-id      : 1
Forward Class  : be                      Profile      : Out
Ring Queue    : 1                      Ring Profile : Out
Scope         : Template                Policy Type  : Default
Description    : Default network QoS policy.
=====
QoS Network Policy
=====
-----
Network Policy (4)
Policy-id      : 4
Ring Queue    : 1                      Ring Profile : Out
Scope         : Template                Policy Type  : Ring
Description    : (Not Specified)
=====
QoS Network Policy
=====
-----
Network Policy (111)
Policy-id      : 111
Forward Class  : be                      Profile      : Out
Scope         : Template                Policy Type  : IpInterface
Description    : (Not Specified)
=====

```

Table 35: Ethernet ring network policy field descriptions

Label	Description
Policy-id	The ID that uniquely identifies the policy
Forward Class	The forwarding class name
Profile	Whether the DSCP mapping pertains to in-profile or out-of-profile traffic
Ring Queue	The queue assigned to the policy
Ring Profile	The profile assigned to the policy: In or Out
Scope	The policy scope: Exclusive or Template
Policy Type	The type of policy: Ip-interface, Ring, or Default

Label	Description
Description	A text string that helps identify the policy's context in the configuration file

sgt-qos

Syntax

sgt-qos

Context

show>router

Description

This command displays QoS information about self-generated traffic.

application

Syntax

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

Context

show>router>sgt-qos

Description

This command displays application QoS settings.

Parameters

app-name

the specified application

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



Note:

- PTP in the context of SGT QoS is defined as Precision Timing Protocol and is an application in the 7705 SAR. The PTP application name is also used in areas such as event-control and logging. Precision Timing Protocol is defined in IEEE 1588-2008.
- PTP in the context of IP filters is defined as Performance Transparency Protocol. IP protocols can be used as

IP filter match criteria; the match is made on the 8-bit protocol field in the IP header.

dscp

specifies to show all DSCP applications

dot1p

specifies to show all dot1p applications

Output

The following output is an example of application QoS information, and [Table 36: Application QoS field descriptions](#) describes the fields.

Output example

```
A:7705:Dut-B# show router sgt-qos application
=====
DSCP Application Values
=====
Application          Configured DSCP Value      Default DSCP Value(s)
-----
bgp                   none                       nc1
cflowd                none                       nc1
dhcp                  none                       nc1, af41, nc2
dns                   none                       af41
ftp                   none                       af41
icmp                  none                       be, nc1
igmp                  none                       nc1
ldp                   none                       nc1
mcfw                  none                       nc1
mld                   none                       nc1
msdp                  none                       nc1
ndis                  none                       nc1, nc2
ntp                   none                       nc1
ospf                  none                       nc1
pcep                  none                       nc1
pim                   none                       nc1
ptp                   none                       nc1
radius                none                       nc1
rip                   none                       nc1
rsvp                  none                       nc1
snmp                  none                       af41
snmp-notification    none                       af41
ssh                   none                       af41
syslog                none                       af41
tacplus               none                       af41
telnet                none                       af41
tftp                  none                       af41
traceroute            none                       be
vrrp                  none                       nc1
=====

Dot1p Application Values
=====
Application          Configured Dot1p Value      Default Dot1p Value
-----
arp                   none                       7
isis                  none                       7
=====
A:7705:Dut-B#
```

Table 36: Application QoS field descriptions

Label	Description
Application	The DSCP or dot1p application
Configured DSCP Value	The DSCP name or value assigned to the application; if you assign a value to the application (0 to 63), the DSCP name that maps to the value is displayed
Default DSCP Value(s)	The default DSCP value Some applications have multiple DSCP default values depending on the context or service
Configured Dot1p Value	The dot1p priority assigned to the application (applies only to ARP and IS-IS)
Default Dot1p Value	The default dot1p value

dscp-map

Syntax

dscp-map [*dscp-name*]

Context

show>router>sgt-qos

Description

This command displays the DSCP-to-FC mappings.

Parameters

dscp-name

the specified DSCP name. [Table 31: Valid DSCP names](#) lists the valid DSCP names.

Output

The following output is an example of DSCP-to-FC mapping information, and [Table 37: DSCP-to-FC mapping field descriptions](#) describes the fields.

Output example

```
A:ALU-1# show router sgt-qos dscp-map
=====
DSCP to FC Mappings
=====
DSCP Value      FC Value      Default FC Value
-----
be              nc
cp1            be              be
```


cp2	be	be
cp3	be	be
cp4	be	be
cp5	be	be
cp6	be	be
cp7	be	be
cs1	be	be
cp9	be	be
af11	af	af
cp11	be	be
af12	af	af
cp13	be	be
af13	af	af
cp15	be	be
cs2	be	be
cp17	be	be
af21	l1	l1
cp19	be	be
af22	l1	l1
cp21	be	be
af23	l1	l1
cp23	be	be
cs3	be	be
cp25	be	be
af31	l1	l1
cp27	be	be
af32	l1	l1
cp29	be	be
af33	l1	l1
cp31	be	be
cs4	be	be
cp33	be	be
af41	nc	nc
cp35	be	be
af42	af	h2
cp37	be	be
af43	h2	h2
cp39	be	be
cs5	be	be
cp41	be	be
cp42	be	be
cp43	be	be
cp44	be	be
cp45	be	be
ef	ef	ef
cp47	be	be
nc1	nc	nc
cp49	be	be
cp50	h2	h2
cp51	be	be
cp52	be	be
cp53	be	be
cp54	be	be
cp55	be	be
nc2	nc	nc
cp57	be	be
cp58	be	be
cp59	be	be
cp60	be	be
cp61	be	be
cp62	be	be
cp63	be	be
=====		
A:ALU-1#		

Table 37: DSCP-to-FC mapping field descriptions

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

5 Network queue QoS policies

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

Topics in this chapter include:

- [Overview](#)
- [Basic configuration](#)
- [Service management tasks](#)
- [Network queue QoS policy command reference](#)

5.1 Overview

Network queue policies define the network queuing characteristics on the network adapter cards.

There is one default network queue policy. Each policy can have up to 16 ingress queues (8 unicast and 8 multipoint). The default policies cannot be deleted but can be copied and the copy can be modified. The default policies are identified as **network-queue default**.

Default network queue policies are applied to adapter card network ingress ports at the adapter card level, and to network egress ports at the port level. You must explicitly create and then associate other network queue QoS policies.

5.2 Basic configuration

This section contains the following topics related to creating and applying network queue QoS policies:

- [Configuring a network queue QoS policy](#)
- [Creating a network queue QoS policy](#)
- [Applying network queue QoS policies](#)
- [Configuring per-VLAN network egress shapers](#)
- [Configuring a CIR for network egress unshaped VLANs](#)
- [Default network queue QoS policy values](#)

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

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

5.2.1 Configuring 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 network ingress and network egress ports, as well as the ports on a ring adapter card, which includes the ring ports, the add/drop port, and the v-port. See [Default network queue QoS policy values](#) for default values.

Perform the following when creating a network queue policy:

- Enter a network queue policy name. The system does not dynamically assign a name.
- Include a description. The description provides a brief overview of policy features.
- Assign a forwarding class. You can assign a forwarding class to a specific queue after the queue has been created.

5.2.2 Creating a network queue QoS policy

By default, all network queue policies are created with queue 1 and (multipoint) queue 9 applied to the policy. Similarly, when an FC is created within a network queue policy, the default unicast queue 1 and multicast queue 9 are assigned to the FC. The **multipoint** keyword applies to queues 9 to 16.

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

CLI syntax:

```
config>qos
  network-queue policy-name
    description description-string
    fc fc-name
      multicast-queue queue-id
      queue queue-id
    queue queue-id [multipoint] [queue-type]
      adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
      cbs percent
      high-prio-only percent
      mbs percent
      rate percent [cir percent]
      slope-policy name
```

The following example creates a unicast and a multipoint network queue policy.

Example:

```
ALU-1# config>qos# network-queue NQ1 create
config>qos>network-queue$ description "NetQueue1"
config>qos>network-queue$ fc be create
config>qos>network-queue>fc$ exit
config>qos>network-queue# queue 10 multipoint create
config>qos>network-queue>queue# exit
config>qos>network-queue$ fc h1 create
config>qos>network-queue>fc$ multicast-queue 10
config>qos>network-queue>fc$ exit
config>qos>network-queue# exit
config>qos# exit
ALU-1#
```

The following output displays the configuration for NQ1.

```

ALU-1>config>qos# network-queue NQ1
ALU-1>config>qos>network-queue# info detail
-----
      description "NetQueue1"
      queue 1 auto-expedite create
        no avg-frame-overhead
        rate 100 cir 0
        adaptation-rule pir closest cir closest
        mbs 5
        cbs 0.10
        high-prio-only 10
        slope-policy "default"
      exit
      queue 9 multipoint auto-expedite create
        no avg-frame-overhead
        rate 100 cir 0
        adaptation-rule pir closest cir closest
        mbs 5
        cbs 0.10
        high-prio-only 10
        slope-policy "default"
      exit
      queue 10 multipoint auto-expedite create
        no avg-frame-overhead
        rate 100 cir 0
        adaptation-rule pir closest cir closest
        mbs 5
        cbs 0.10
        high-prio-only 10
        slope-policy "default"
      exit
      fc be create
        multicast-queue 9
        queue 1
      exit
      fc hl create
        multicast-queue 10
        queue 1
      exit
-----
ALU-1>config>qos>network-queue#

```

5.2.3 Applying network queue QoS policies

Apply network queue policies to the following entities:

- [Adapter cards](#)
- [Network ports](#)

5.2.3.1 Adapter cards

Use the following CLI syntax to apply a network queue policy to any adapter card network ingress port or v-port on a ring adapter card. Use the **ring** command to apply a network queue policy to a ring add/drop port. You cannot assign a network queue policy to an add/drop port while the policy is referenced by a non-ring port, and vice versa.

The first example applies a network queue policy to any network ingress port or v-port on a ring adapter card. The second example applies to network queue policy to an add/drop port.

CLI syntax:

```
config>card
      mda mda-slot
        network
          ingress
            queue-policy name
      ring
        add-drop-port-queue-policy name
```

Example:

```
configure card 1
config>card# mda 1
config>card>mda# network
config>card>mda>network# ingress
config>card>mda>network>ingress# queue-policy NQ1
config>card>mda>network>ingress# exit
config>card>mda>network# exit
config>card>mda# exit
config>card# exit
```

The following output displays network ingress queue policy NQ1 applied to the adapter cards.

```
A:ALU-1# configure card 1
*A:ALU-1>config>card# info
-----
      card-type iom-sar
      mda 1
        mda-type a16-chdslv2
        network
          ingress
            queue-policy "NQ1"
          exit
        exit
      exit
    mda 2
      mda-type a8-lgb-v3-sfp
      network
        ingress
          queue-policy "NQ1"
        exit
      exit
    exit
  mda 3
    mda-type a8-lgb-v3-sfp
    network
      ingress
        queue-policy "NQ1"
      exit
    exit
  exit
mda 4
  mda-type a16-chdslv2
  network
    ingress
      queue-policy "NQ1"
    exit
  exit
exit
-----
```

```
*A:ALU-1>config>card#
```

Use the following CLI syntax to apply a network queue policy to the add/drop port. A network queue policy for an add/drop port applies to traffic flowing from the bridging domain to the IP domain (that is, from the add/drop port to the v-port).

Example:

```
config# card 1
config>card# mda 8
config>card>mda# network ring
config>card>mda>network>ring# add-drop-port-queue-policy "adp_queue_
policy"
```

The following output displays the configuration for a network queue policy applied to the add/drop port.

```
*A:7705custDoc:Sar18>config>card>mda>network# info detail
-----
        ingress
            fabric-policy 1
            queue-policy "default"
        exit
        ring
            qos-policy 5
            add-drop-port-queue-policy "adp_queue_policy"
        exit
-----
*A:7705custDoc:Sar18>config>card>mda>network#
```

5.2.3.2 Network ports

Use the following CLI syntax to apply network queue policy NQ1 to a network port.

CLI syntax:

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

Example:

```
ALU-1# config# port 1/1/1
config>port# ethernet
config>port>ethernet# network
config>port>ethernet>network# queue-policy NQ1
config>port>ethernet>network# exit
ALU-1#
```

The following output displays a network port configuration.

```
*A:ALU-1>config>port# info
-----
        ethernet
            network
                queue-policy "NQ1"
            exit
        exit
        no shutdown
-----
```

```
*A:ALU-1>config>port#
```

Use the following CLI syntax to apply a network queue policy to v-port egress traffic.

A network queue policy for v-port egress traffic applies DSCP and LSP EXP classification to traffic flowing from the switching fabric in the IP domain toward the add/drop port in the bridging domain.

Example:

```
config# port 1/8/v-port
config>port# ethernet network
config>port>ethernet>network# queue-policy "egr_vp_q_policy"
```

The following output displays the configuration for an egress queue policy applied to the v-port.

```
*A:7705custDoc:Sar18>config>port>ethernet>network# info detail
-----
queue-policy "egr_vp_q_policy"
scheduler-mode 16-priority
-----
```

A network queue policy for a ring port must be configured as a ring type using the **network-policy-type** keyword. The policy operates on ring port egress traffic.

Use the following CLI syntax to apply a network queue policy to a ring port:

CLI syntax:

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

Example:

```
config# port 1/8/1
config>port# ethernet network
config>port>ethernet>network# queue-policy "NQpolicy5"
```

The following output displays the configuration for a network queue policy applied to a ring port.

```
*A:7705custDoc:Sar18>config>port>ethernet>network# info detail
-----
queue-policy "NQpolicy5"
scheduler-mode 16-priority
-----
```

5.2.4 Configuring per-VLAN network egress shapers

Per-VLAN network egress shapers can be configured for network interfaces. See the 7705 SAR Router Configuration Guide for command descriptions.

The **queue-policy** command is used to enable and disable the network egress per-VLAN shapers on a per-interface basis. If the **no queue-policy** command is used, the VLAN (that is, the interface) defaults to unshaped mode. The **agg-rate-limit** command cannot be accessed unless a network queue policy is assigned to the interface.

Use the following CLI syntax to configure a per-VLAN network egress shaper on a network interface:

CLI syntax:

```
config>router>interface#
    egress
      queue-policy name
      agg-rate-limit agg-rate [cir cir-rate]
```

5.2.5 Configuring a CIR for network egress unshaped VLANs

To provide arbitration between the bulk (aggregate) of unshaped VLANs and the shaped VLANs, assign a rate to the unshaped VLANs. See the 7705 SAR Interface Configuration Guide for command descriptions.

Use the following CLI syntax to configure a CIR for the bulk of network egress unshaped VLANs:

CLI syntax:

```
config>port>ethernet#
    network
      egress
        unshaped-if-cir cir-rate
```

5.2.6 Default network queue QoS policy values

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

Table 38: Default network queue policy definitions

Forwarding class	Queue	Definition	Queue	Definition
Network-Control (nc)	8	Rate = 100% CIR = 10% MBS = 2.5% CBS = 0.25% High-Prio-Only = 10%	16	Rate = 100% CIR = 10% MBS = 2.5% CBS = 0.1% High-Prio-Only = 10%
High-1 (h1)	7	Rate = 100% CIR = 10% MBS = 2.5% CBS = 0.25% High-Prio-Only = 10%	15	Rate = 100% CIR = 10% MBS = 2.5% CBS = 0.1% High-Prio-Only = 10%
Expedited (ef)	6	Rate = 100% CIR = 100% MBS = 5% CBS = 0.75%	14	Rate = 100% CIR = 100% MBS = 5% CBS = 0.1%

Forwarding class	Queue	Definition	Queue	Definition
		High-Prio-Only = 10%		High-Prio-Only = 10%
High-2 (h2)	5	Rate = 100% CIR = 100% MBS = 5% CBS = 0.75% High-Prio-Only = 10%	13	Rate = 100% CIR = 100% MBS = 5% CBS = 0.1% High-Prio-Only = 10%
Low-1 (l1)	4	Rate = 100% CIR = 25% MBS = 2.5% CBS = 0.25% High-Prio-Only = 10%	12	Rate = 100% CIR = 5% MBS = 2.5% CBS = 0.25% High-Prio-Only = 10%
Assured (af)	3	Rate = 100% CIR = 25% MBS = 5% CBS = 0.75% High-Prio-Only = 10%	11	Rate = 100% CIR = 5% MBS = 5% CBS = 0.1% High-Prio-Only = 10%
Low-2 (l2)	2	Rate = 100% CIR = 25% MBS = 5% CBS = 0.25% High-Prio-Only = 10%	10	Rate = 100% CIR = 5% MBS = 5% CBS = 0.1% High-Prio-Only = 10%
Best Effort (be)	1	Rate = 100% CIR = 0% MBS = 5% CBS = 0.1% High-Prio-Only = 10%	9	Rate = 100% CIR = 0% MBS = 5% CBS = 0.1% High-Prio-Only = 10%

The following output displays the network queue policy default configuration.

```

ALU-1>config>qos>network-queue# info detail
-----
description "Default network queue QoS policy."
queue 1 auto-expedite create
no avg-frame-overhead
rate 100 cir 0
adaptation-rule pir closest cir closest
mbs 5
cbs 0.10
high-prio-only 10

```

```
        slope-policy "default"
    exit
    queue 2 auto-expedite create
        no avg-frame-overhead
        rate 100 cir 25
        adaptation-rule pir closest cir closest
        mbs 5
        cbs 0.25
        high-prio-only 10
        slope-policy "default"
    exit
    queue 3 auto-expedite create
        no avg-frame-overhead
        rate 100 cir 25
        adaptation-rule pir closest cir closest
        mbs 5
        cbs 0.75
        high-prio-only 10
        slope-policy "default"
    exit
    queue 4 auto-expedite create
        no avg-frame-overhead
        rate 100 cir 25
        adaptation-rule pir closest cir closest
        mbs 2.5
        cbs 0.25
        high-prio-only 10
        slope-policy "default"
    exit
    queue 5 auto-expedite create
        no avg-frame-overhead
        rate 100 cir 100
        adaptation-rule pir closest cir closest
        mbs 5
        cbs 0.75
        high-prio-only 10
        slope-policy "default"
    exit
    queue 6 auto-expedite create
        no avg-frame-overhead
        rate 100 cir 100
        adaptation-rule pir closest cir closest
        mbs 5
        cbs 0.75
        high-prio-only 10
        slope-policy "default"
    exit
    queue 7 auto-expedite create
        no avg-frame-overhead
        rate 100 cir 10
        adaptation-rule pir closest cir closest
        mbs 2.5
        cbs 0.25
        high-prio-only 10
        slope-policy "default"
    exit
    queue 8 auto-expedite create
        no avg-frame-overhead
        rate 100 cir 10
        adaptation-rule pir closest cir closest
        mbs 2.5
        cbs 0.25
        high-prio-only 10
        slope-policy "default"
```

```
exit
queue 9 multipoint auto-expedite create
  no avg-frame-overhead
  rate 100 cir 0
  adaptation-rule pir closest cir closest
  mbs 5
  cbs 0.10
  high-prio-only 10
  slope-policy "default"
exit
queue 10 multipoint auto-expedite create
  no avg-frame-overhead
  rate 100 cir 5
  adaptation-rule pir closest cir closest
  mbs 5
  cbs 0.10
  high-prio-only 10
  slope-policy "default"
exit
queue 11 multipoint auto-expedite create
  no avg-frame-overhead
  rate 100 cir 5
  adaptation-rule pir closest cir closest
  mbs 5
  cbs 0.10
  high-prio-only 10
  slope-policy "default"
exit
queue 12 multipoint auto-expedite create
  no avg-frame-overhead
  rate 100 cir 5
  adaptation-rule pir closest cir closest
  mbs 2.5
  cbs 0.25
  high-prio-only 10
  slope-policy "default"
exit
queue 13 multipoint auto-expedite create
  no avg-frame-overhead
  rate 100 cir 100
  adaptation-rule pir closest cir closest
  mbs 5
  cbs 0.10
  high-prio-only 10
  slope-policy "default"
exit
queue 14 multipoint auto-expedite create
  no avg-frame-overhead
  rate 100 cir 100
  adaptation-rule pir closest cir closest
  mbs 5
  cbs 0.10
  high-prio-only 10
  slope-policy "default"
exit
queue 15 multipoint auto-expedite create
  no avg-frame-overhead
  rate 100 cir 10
  adaptation-rule pir closest cir closest
  mbs 2.5
  cbs 0.10
  high-prio-only 10
  slope-policy "default"
exit
```

```

queue 16 multipoint auto-expedite create
  no avg-frame-overhead
  rate 100 cir 10
  adaptation-rule pir closest cir closest
  mbs 2.5
  cbs 0.10
  high-prio-only 10
  slope-policy "default"
exit
fc af create
  multicast-queue 11
  queue 3
exit
fc be create
  multicast-queue 9
  queue 1
exit
fc ef create
  multicast-queue 14
  queue 6
exit
fc h1 create
  multicast-queue 15
  queue 7
exit
fc h2 create
  multicast-queue 13
  queue 5
exit
fc l1 create
  multicast-queue 12
  queue 4
exit
fc l2 create
  multicast-queue 10
  queue 2
exit
fc nc create
  multicast-queue 16
  queue 8

```

5.3 Service management tasks

This section describes the following service management tasks:

- [Deleting QoS policies](#)
- [Copying and overwriting QoS policies](#)
- [Editing QoS policies](#)

5.3.1 Deleting QoS policies

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

Use the following CLI syntax to delete a network queue policy.

CLI syntax:

```
config>qos# no network-queue policy-name
```

Example:

```
config>qos# no network-queue NQ1
```

5.3.2 Copying and overwriting QoS policies

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

Use the following CLI syntax to overwrite an existing network queue policy.

CLI syntax:

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

Example:

```
A:ALU-1>config>qos# copy network-queue NQ1 NQ2 overwrite
config>qos# exit
*A:ALU-1#
```

The following output displays the copied policies:

```
*A:ALU-1>config>qos# info
-----
#-----
echo "QoS Policy Configuration"
#-----
    network-queue "NQ1" create
        description "NetQueue1"
        queue 1 create
            rate 10
            mbs 5
            cbs 0.10
            high-prio-only 10
        exit
    queue 9 multipoint create
        rate 10
        mbs 5
        cbs 0.10
        high-prio-only 10
    exit
    fc be create
        multicast-queue 9
        queue 1
    exit
exit
network-queue "NQ2" create
    description "NetQueue1"
    queue 1 create
        rate 10
        mbs 5
        cbs 0.10
        high-prio-only 10
```

```

    exit
    queue 9 multipoint create
        rate 10
        mbs 5
        cbs 0.10
        high-prio-only 10
    exit
    fc be create
        multicast-queue 9
        queue 1
    exit
exit
network-queue "nq1" create
description "NetQ1"
queue 1 create
    rate 10
    mbs 5
    cbs 0.10
    high-prio-only 10
exit
queue 9 multipoint create
    rate 10
    mbs 5
    cbs 0.10
    high-prio-only 10
exit
fc be create
    multicast-queue 9
    queue 1
exit
exit
network-queue "nq3" create
description "NetQ3"
queue 1 create
    mbs 5
    cbs 0.10
    high-prio-only 10
exit
queue 9 multipoint create
    rate 10
    mbs 5
    cbs 0.10
    high-prio-only 10
exit
fc be create
    multicast-queue 9
    queue 1
exit
exit
network-queue "netq2" create
Press any key to continue (Q to quit)

```

5.3.3 Editing QoS policies

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

5.4 Network queue QoS policy command reference

5.4.1 Command hierarchies

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

5.4.1.1 Configuration commands

```

config
- qos
  - [no] network-queue policy-name [create]
    - description description-string
    - no description
    - [no] fc fc-name [create]
      - multicast-queue queue-id
      - no multicast-queue
      - queue queue-id
      - no queue
    - packet-byte-offset [add bytes | subtract bytes | none]
    - no packet-byte-offset
    - queue queue-id [multipoint] [queue-type] [create]
    - no queue queue-id
      - avg-frame-overhead percent
      - no avg-frame-overhead
      - adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
      - no adaptation-rule
      - cbs percent
      - no cbs
      - high-prio-only percent
      - no high-prio-only
      - mbs percent
      - no mbs
      - packet-byte-offset [add bytes | subtract bytes | none]
      - no packet-byte-offset
      - rate percent [cir percent]
      - no rate
      - slope-policy name
      - no slope-policy

```

```

config
- [no] port port-id
  - ethernet
    - egress-rate sub-rate [include-fcs] [allow-eth-bn-rate-changes] [hold-time hold-
time]
    - no egress-rate
  - network
    - scheduler-mode {16-priority}

```


5.4.1.2 Operational commands

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

5.4.1.3 Show commands

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

5.4.2 Command descriptions

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

5.4.2.1 Configuration commands

- [Generic commands](#)
- [Network queue QoS policy commands](#)
- [Network queue QoS policy forwarding class commands](#)

5.4.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>network-queue

config>qos>network

config>qos>sap-egress

config>qos>sap-ingress

Description

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

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

Default

n/a

Parameters

description-string

a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

5.4.2.1.2 Network queue QoS policy commands

network-queue

Syntax

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

Context

config>qos

Description

This command creates a context to configure a network queue policy. Network queue policies define the ingress and egress network queuing at the adapter card network node level.

Network queue policies define ingress and egress network queues similar to a service ingress QoS policy.

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

the name of the network queue policy

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

create

keyword used to create a network queue policy

packet-byte-offset

Syntax

packet-byte-offset [add *bytes* | subtract *bytes* | none]

no packet-byte-offset

Context

config>qos>network-queue

config>qos>network-queue>queue

Description

This command is used to modify the size of the packet that schedulers operate on. Modification only impacts schedulers and queue statistics. The actual packet size is not modified, nor can it be. Only the size used by the schedulers to determine the scheduling is changed. The **packet-byte-offset** command is meant to be a mechanism that can be used to compensate for downstream encapsulation or header removal. The scheduling rates are affected by the offset, as well as the statistics (accounting) associated with the queue. The **packet-byte-offset** command does not affect port-level and service-level statistics. It only affects the queue statistics. The network-queue policy applies in both the ingress and egress directions.

The **add** and **subtract** keywords are mutually exclusive. Either **add**, **subtract**, or **none** must be specified.

There are three modes of **packet-byte-offset** operation:

- **no packet-byte-offset** – enables legacy behavior so that no modification is performed
- **packet-byte-offset** – automatic adjustment mode. Rates apply to packets based on the received packet size at ingress (this is also known as packet size on the wire, less the Layer 1 headers, the inter-frame GAP and the Preamble) and to the transmitted packet size at egress, which includes 4 bytes of Ethernet FCS. At ingress, all internal headers and associated service headers are discounted during scheduling operation. At egress, 4 bytes are added to accommodate for Ethernet FCS.
- **packet-byte-offset [add bytes | subtract bytes]** – automatic correction followed by addition or subtraction of a specified number of bytes. This command first performs the **packet-byte-offset** operation as captured above and then adds or subtracts a certain number of bytes. Rates apply to packets based on the size of the packet at the ingress or egress port plus or minus an offset.

Packet byte offset configuration can be applied at the policy level, in which case it applies to all of the queues within the policy, or at the individual queue level so that it applies only to a specific queue.

The **no** version of this command enables legacy 7705 SAR behavior where the queue rates are relative to the packet size with the internal fabric header added, but without the FCS.

Parameters

add bytes

after automatic adjustment for internal headers (for example, added FCS or removal of internal service/overhead), adds the specified number of bytes to each packet associated with the queue for scheduling and accounting purposes. From the queue's perspective, the packet size is increased by the amount being added to each packet.

Values 2 to 62, in steps of 2

subtract bytes

after automatic adjustment for internal headers (for example, added FCS or removal of internal service/overhead), subtracts the specified number of bytes from each packet associated with the queue for scheduling and accounting purposes. From the queue's perspective, the packet size is reduced by the amount being subtracted from each packet.

Values 2 to 62, in steps of 2

none

the packet size is left unchanged

queue

Syntax

queue *queue-id* **multipoint** [*queue-type*] [**create**]
no queue *queue-id*

Context

config>qos>network-queue

Description

This command enables the context to configure a QoS network-queue policy queue. Network queues are created with default queue 1 (non-multipoint) and queue 9 (multipoint) automatically assigned.

The **queue** command with the **multipoint** keyword allows the creation of multipoint queues. Only multipoint queues can receive ingress packets that need flooding to multiple destinations. By separating the unicast traffic from multipoint traffic at network ingress and handling the traffic on separate multipoint queues, special handling of the multipoint traffic is possible. Each queue acts as an accounting and (optionally) shaping device, offering precise control over potentially expensive broadcast, multicast, and unknown unicast traffic. Only the back-end support of multipoint traffic (between the forwarding class and the queue based on forwarding type) needs to be defined. The individual classification rules used to place traffic into forwarding classes are not affected. Queues must be defined as multipoint at the time of creation within the policy.

The multipoint queues are for multipoint-destined service traffic. Within non-multipoint services, such as Epipe services, all traffic is considered unicast 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 queue.

The **no** form of this command removes the forwarding class-to-queue mapping, causing the forwarding class to use the default queue instead. When a queue is removed, any pending accounting information for each network queue created because of the definition of the queue in the policy is discarded.

Parameters

<i>queue-id</i>	the queue identifier for the queue, expressed as an integer. The <i>queue-id</i> uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.
Values	1 to 8 (unicast) 9 to 16 (multipoint)
Default	1 (unicast) 9 (multipoint)

multipoint

specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be explicitly mapped to the forwarding class broadcast, multicast, or unknown unicast (BMU) ingress traffic. If you attempt to map forwarding class unicast traffic to a queue

designated as multipoint, an error is generated and no changes are made to the current unicast traffic queue mapping.

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

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

Values multipoint or not present

Default not present (the queue is created as a unicast queue)

queue-type

the **expedite**, **best-effort** and **auto-expedite** queue types are mutually exclusive. Each defines the method that the system uses to service the queue from a hardware perspective. A keyword must be specified at the time the queue is created in the network-queue policy. If an attempt is made to change the keyword after the queue is initially defined, an error is generated.

expedite

the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue

best-effort

the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue

auto-expedite

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

Values expedite, best-effort, auto-expedite

Default auto-expedite

create

keyword used to create a network QoS policy

avg-frame-overhead

Syntax

avg-frame-overhead *percent*

no avg-frame-overhead

Context

config>qos>network-queue>queue

Description

The 7705 SAR does not support the **avg-frame-overhead** command. It is always set to 0 and cannot be modified.

adaptation-rule

Syntax

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

no adaptation-rule

Context

config>qos>network-queue>queue

Description

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

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

Default

adaptation-rule pir closest cir closest

Parameters

pir

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

cir

defines the constraints enforced when adapting the CIR rate defined within the **queue queue-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule

specifies the adaptation rule to be used while computing the operational CIR or PIR value. The **max** (maximum), **min** (minimum), and **closest** parameters are mutually exclusive.

Values **max** – causes the network processor to be programmed at an operational rate that is less than the configured PIR or CIR rate by up

to 1.0%. For a network processor on a Gen-3 adapter card or platform, the average difference between the operational and the configured CIR rate is 2.0% (for frame sizes less than 2049 bytes) or 4.0% (for other frame sizes).

min – causes the network processor to be programmed at an operational rate that is greater than the configured PIR or CIR rate by up to 1.0%. For a network processor on a Gen-3 adapter card or platform, the average difference between the operational and the configured CIR rate is 2.0% (for frame sizes less than 2049 bytes) or 4.0% (for other frame sizes).

closest – causes the network processor to be programmed at an operational rate that is closest to the configured PIR or CIR rate

cbs

Syntax

cbs *percent*

no cbs

Context

config>qos>network-queue>queue

Description

This command specifies the relative amount of reserved buffers for a specific ingress network adapter card forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage. The resultant CBS size can be larger than the MBS. This will result in a portion of the CBS for the queue to be unused and therefore should be avoided.

The **no** form of this command returns the CBS size for the queue to the default for the forwarding class.

Default

The following table lists the **cbs** forwarding class defaults.

Table 39: CBS forwarding class defaults

Forwarding class	Forwarding class label	Unicast queues		Multicast queues	
		Queue ID	Default CBS (%)	Queue ID	Default CBS (%)
Network-Control	nc	8	0.25	16	0.1
High-1	h1	7	0.25	15	0.1
Expedited	ef	6	0.75	14	0.1
High-2	h2	5	0.75	13	0.1
Low-1	l1	4	0.25	12	0.1

Forwarding class	Forwarding class label	Unicast queues		Multicast queues	
		Queue ID	Default CBS (%)	Queue ID	Default CBS (%)
Assured	af	3	0.75	11	0.1
Low-2	l2	2	0.25	10	0.1
Best-Effort	be	1	0.1	9	0.1

Special cases

Forwarding class queue on egress network ports or bundles

the total reserved buffers based on the total percentages can exceed 100%. This might not be desirable and should be avoided as a rule of thumb. If the total percentage equals or exceeds 100% of the queue size, no buffers will be available in the shared portion of the pool. Any queue exceeding its CBS size will experience a hard drop on all packets until it drains below this threshold.

Forwarding class queue on ingress adapter cards

the total reserved buffers based on the total percentages can exceed 100%. This might not be desirable and should be avoided as a rule of thumb. If the total percentage equals or exceeds 100% of the queue size, no buffers will be available in the shared portion of the pool. Any queue exceeding its CBS size will experience a hard drop on all packets until it drains below this threshold.

Parameters

percent

the percent of buffers reserved from the total queue space, expressed as a decimal integer. If 10 MB is the total buffer value in the queue, a value of 10 would reserve 1 MB (10%) of buffer space for the forwarding class queue. The value 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can be applied for scheduling purposes).

Values 0.00 to 100.00

high-prio-only

Syntax

high-prio-only *percent*

no high-prio-only

Context

config>qos>network-queue>queue

Description

The **high-prio-only** command allows the reservation of queue buffers for use exclusively by in-profile packets as a default condition for access buffer queues for this network queue policy. For network queues, in-profile packets are high priority, and out-of-profile packets are low priority.



Note: When a low-priority RED/WRED slope is enabled on a queue, the high-prio-only setting is not used. When that slope is disabled, then the high-prio-setting is used.

Modifying the current MBS for the queue through the **mbs** command will cause the default **high-prio-only** function to be recalculated and applied to the queue.

The **no** form of this command restores the default value.

Default

The following table lists the **high-prio-only** forwarding class defaults.

Table 40: High-prio-only forwarding class defaults

Forwarding class	Forwarding class label	Unicast queues		Multicast queues	
		Queue ID	Default high-prio-only	Queue ID	Default high-prio-only
Network-Control	nc	8	10	16	10
High-1	h1	7	10	15	10
Expedited	ef	6	10	14	10
High-2	h2	5	10	13	10
Low-1	l1	4	10	12	10
Assured	af	3	10	11	10
Low-2	l2	2	10	10	10
Best-Effort	be	1	10	9	10

Parameters

percent

the amount of queue buffer space reserved for in-profile packets, expressed as a decimal percentage

Values 0 to 100 | default

mbs

Syntax

mbs *percent*

no mbs**Context**

```
config>qos>network-queue>queue
```

Description

This command specifies the relative amount of the queue space for the maximum buffers for a specific ingress network adapter card forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.

The maximum burst size (MBS) value is used by a queue to determine whether it has exhausted its total allowed buffers while enqueueing packets. Once the queue has exceeded its maximum amount of buffers, all packets are discarded until the queue transmits a packet. A queue that has not exceeded its MBS is not guaranteed that a buffer will be available when needed or that the packet's RED/WRED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED/WRED slope parameters for the needs of the network queues.

The MBS can sometimes be smaller than the CBS. This will result in a portion of the CBS for the queue to be unused and should be avoided.

The **no** form of this command returns the MBS for the queue to the default for the forwarding class.

Default

The following table lists the **mbs** forwarding class defaults.

Table 41: MBS forwarding class defaults

Forwarding class	Forwarding class label	Unicast queue		Multicast queue	
		Queue ID	Default MBS	Queue ID	Default MBS
Network-Control	nc	8	2.5	16	2.5
High-1	h1	7	2.5	15	2.5
Expedited	ef	6	5	14	5
High-2	h2	5	5	13	5
Low-1	l1	4	2.5	12	2.5
Assured	af	3	5	11	5
Low-2	l2	2	5	10	5
Best-Effort	be	1	5	9	5

Special cases**Forwarding class queue on egress network ports or bundles**

the total MBS settings for all network egress queues on the port or channel based on the total percentages can exceed 100%. Some oversubscription can be desirable to allow

exceptionally busy forwarding classes more access to buffer space. The proper use of CBS settings will ensure that oversubscribing MBS settings will not starve other queues of buffers when needed.

Forwarding class queue on ingress adapter cards

the **mbs** value is used to calculate the queue's MBS based on the total amount of buffer space allocated to the network ingress queue on the adapter card.

The total MBS settings for all network egress queues on the port or channel based on the total percentages can exceed 100%. Some oversubscription can be desirable to allow exceptionally busy forwarding classes more access to buffer space. The correct use of CBS settings will ensure that oversubscribing MBS settings will not starve other queues of buffers when needed.

Parameters

percent

the percent of buffers from the total queue space allowed for the maximum amount of buffers, expressed as a decimal percentage

Values 0.00 to 100.00

rate

Syntax

rate *percent* [**cir** *percent*]

no rate

Context

config>qos>network-queue>queue

Description

This command defines the administrative peak information rate (PIR) and the administrative committed information rate (CIR) parameters for the queue. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the percentage at which the system prioritizes the queue over other queues competing for the same bandwidth.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP ingress or SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (100, 0).

Parameters

percent

defines the percentage of the maximum rate allowed for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate**

command has not been executed, the default PIR of 100 is assumed. Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned. The PIR rate has a minimum value of 8 kb/s for all hardware.

Values1 to 100

Default100


cir percent

defines the percentage of the maximum rate allowed for the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer.

The CIR rate has a minimum value of 8 kb/s.

Values0 to 100

Default0

 **Note:** If a specified percentage results in a PIR or CIR that is lower than the minimum rate, the system rounds up the CIR or PIR to the minimum rate.

slope-policy

Syntax

slope-policy *name*

no slope-policy

Context

config>qos>network-queue>queue

Description

This command specifies the name of slope policy associated with the network queue.

Parameters

name

specifies the name for the slope policy

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

Defaultdefault

egress-rate

Syntax

egress-rate *sub-rate* [**include-fcs**] [**allow-eth-bn-rate-changes**] [**hold-time** *hold-time*]

no egress-rate

Context

config>port>ethernet

Description

This command configures the rate of traffic leaving the network.

With the **include-fcs** option, the egress rate limit is applied to the traffic rate egressing the port with the 4-byte Ethernet FCS field included. If this option is not configured, the egress rate limit is applied to the traffic rate egressing the port without the 4-byte Ethernet FCS field included, and the actual rate of packets leaving the port is slightly higher than the configured egress rate value.

The **include-fcs** option is not supported on the 7705 SAR-A Fast Ethernet ports (ports 9 to 12) or 4-port SAR-H Fast Ethernet module. On the 6-port SAR-M Ethernet module, the **include-fcs** option is always on and cannot be disabled to compensate for the 4-byte FCS.

The **allow-eth-bn-rate-changes** option enables the Y.1731 ETH-BN client MEP option on the port. In applications such as a point-to-point microwave link, where degradation on the line can result in reduced link bandwidth, the egress rate can be dynamically changed based on the available bandwidth on the link as indicated by the ETH-BN server. When enabled, the received rate overrides the configured sub-rate for the port. For information about ETH-BN, including which Ethernet ports support this functionality, see the 7705 SAR OAM and Diagnostics Guide, "ITU-T Y.1731 Ethernet Bandwidth Notification (ETH-BN)".

The bandwidth indicated by the ETH-BN server includes the FCS; therefore, the **include-fcs** option must be selected if the **allow-eth-bn-rate-changes** option is selected or the dynamically changed bandwidth will not match the intended rate.

The *hold-time* is used to limit the number of bandwidth changes as requested by the ETH-BN server. After a rate change occurs based on a Bandwidth Notification Message (BNM), any BMN received before the hold timer expires will be ignored.

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

Default

no egress-rate

Parameters

sub-rate

the egress rate in kb/s

Values 1 to 10000000

include-fcs

the egress rate limit is applied to the traffic rate egressing the port with the 4-byte Ethernet FCS field included. This option must be selected if the allow-eth-bn-rate-changes option is selected; otherwise, the dynamically changed bandwidth will not match the intended rate.

allow-eth-bn-rate-changes

enables the Y.1731 ETH-BN client MEP option on the port. The egress rate will be dynamically changed to the bandwidth indicated in messages received from an ETH-BN server MEP. When enabled, the received rate overrides the configured sub-rate for the port.

hold-time

configures the hold time for egress rate bandwidth changes based on a received BNM, in seconds

Values 1 to 600

Default 5

scheduler-mode

Syntax

scheduler-mode {16-priority}

Context

config>port>ethernet>network

Description

This command selects the network-side scheduling option for Ethernet ports on the equipment listed in [Table 10: Scheduling modes supported by adapter cards and ports at network egress](#) and [Table 12: Scheduling modes supported by adapter cards and ports at network ingress](#).

On the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X, **scheduler-mode** is permanently set to support 4-priority and is not user-configurable. On all other Ethernet Adapter cards, modules, and platforms listed in Table 10 and Table 12, **scheduler-mode** can only be configured to 16-priority.

Default

16-priority—8-port Gigabit Ethernet Adapter card, 10-port 1GigE/1-port 10GigE X-Adapter card, 2-port 10GigE (Ethernet) Adapter card, 2-port 10GigE (Ethernet) module, 4-port SAR-H Fast Ethernet module, 6-port SAR-M Ethernet module, Packet Microwave Adapter card, 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M, and 7705 SAR-Wx Ethernet ports (cannot be changed)

Parameters**16-priority**

sets the 16-priority scheduling option for the cards, modules, and platforms listed under Default

5.4.2.1.3 Network queue QoS policy forwarding class commands

fc

Syntax

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

Context

config>qos>network-queue

Description

The **fc** is created with default unicast *queue-id* 1 and default multicast *queue-id* 9 automatically configured. The specified queues contain the PIR, CIR, CBS, and MBS configurations.

Use the [multicast-queue](#) and [queue](#) commands to change the **fc** *queue-id* assignments from their default queue assignments.

The **no** form of this command restores the default queue.

Parameters

fc-name

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

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

create

keyword used to create a forwarding class policy

multicast-queue

Syntax

multicast-queue *queue-id*

no multicast-queue

Context

config>qos>network-queue>fc

Description

This command overrides the default multicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. After the forwarding class mapping is executed, all multicast traffic at network ingress using this policy is forwarded using the *queue-id*. Use the [queue](#) *queue-id* **multipoint** command to create the specified *queue-id*.

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

The **no** form of the command sets the multicast forwarding type *queue-id* back to the default queue (queue 9).

Parameters

<i>queue-id</i>	
	an existing multipoint queue defined in the config>qos>network-queue context.
Values	9 to 16
Default	9

queue

Syntax

queue *queue-id*
no queue

Context

config>qos>network-queue>fc

Description

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

The **no** form of this command removes the *queue-id* from the network-queue policy and from any existing network ingress or network egress ports using the policy, and sets the *queue-id* back to the default queue (queue 1).

Parameters

<i>queue-id</i>	
	the <i>queue identifier</i> for the queue, expressed as an integer. The <i>queue-id</i> uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.
Values	1 to 8
Default	1

5.4.2.2 Operational commands

copy

Syntax

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

Context

config>qos

Description

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

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

Parameters

network-queue *src-name dst-name*

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

overwrite

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

SAR12>config>qos# copy network-queue nq1 nq2

MINOR: CLI Destination "nq2" exists - use {overwrite}.

SAR12>config>qos# copy network-queue nq1 nq2 overwrite

5.4.2.3 Show commands



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

network-queue

Syntax

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

Context

show>qos

Description

This command displays network queue policy information. This includes queue parameters information, forwarding class-to-queue mappings, and network port/adapter card queue associations.

Parameters

network-queue-policy-name
the name of the network queue policy

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

detail
displays detailed network queue information

Output

The following output is an example of network queue policy information, and [Table 42: Network queue policy field descriptions](#) describes the fields.

Output example

```
ALU-1>show>qos# network-queue policy102
=====
QoS Network Queue Policy
=====
-----
Network Queue Policy (policy102)
-----
Policy      : policy102
Description : (Not Specified)
-----
Associations
-----
No Matching Entries
=====

ALU-1>show>qos# network-queue policy102 detail
=====
QoS Network Queue Policy
=====
-----
Network Queue Policy (policy102)
-----
Policy      : policy102
Description : (Not Specified)
Pkt.Byte Offset: sub 62
-----
Queue CIR      PIR      CBS      MBS      HiPrio AvgOvrhd Packet Slope-Policy
      CIR Rule  PIR Rule              Offset
-----
1      0      100      def      def      10      0.00      def      default
      closest closest
```

```

2      0      100      def      def      def      0.00      def      default
      closest closest
9      0      100      def      def      10      0.00      def      default
      closest closest
10     0      100      def      def      def      0.00      def      default
      closest closest
-----
FC      UCastQ      MCastQ
-----
h2      1          9
h1      1          9
-----
Associations
-----
No Matching Entries
=====
ALU-1>show>qos#

```

Table 42: Network queue policy field descriptions

Label	Description
Policy	The policy name that uniquely identifies the policy
Description	A text string that helps identify the policy's context in the configuration file
Pkt.Byte Offset	Indicates the value of the packet byte offset applied to the packet for scheduling, if applicable. A value of "default" indicates that legacy mode packet scheduling is in use, in which packets are scheduled based on size including internal overhead.
Queue	The queue ID
CIR	The committed information rate
CIR Rule	min - the operational CIR for the queue will be equal to or greater than the administrative rate specified using the rate command except where the derived operational CIR is greater than the operational PIR. If the derived operational CIR is greater than the derived operational PIR, the operational CIR will be made equal to the operational PIR. max - the operational CIR for the queue will be equal to or less than the administrative rate specified using the rate command closest - the operational CIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational PIR
PIR	The peak information rate
PIR Rule	min - the operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command

Label	Description
	max - the operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command
	closest - the operational PIR for the queue will be the rate closest to the rate specified using the rate command
CBS	The committed burst size
MBS	The maximum burst size
HiPrio	The high-priority value
AvgOvrhd	The average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire
Packet Offset	The value of the packet byte offset applied to the queue. A value of "default" indicates that legacy mode packet scheduling is in use, in which packets are scheduled based on size including internal overhead.
Slope-Policy	The slope policy for the queue
FC	The value of a predefined forwarding class
UCastQ	The specific unicast queue to be used for packets in the forwarding class
MCastQ	The specific multicast queue to be used for packets in the forwarding class
Associations	The unique service and customer identifiers

6 Service egress and ingress QoS policies

This chapter provides information to configure service egress and ingress QoS policies using the command line interface.

Topics in this chapter include:

- [Overview](#)
- [Basic configuration](#)
- [Service management tasks](#)
- [Service egress and ingress QoS policy command reference](#)

6.1 Overview

There is one default policy for service ingress and one default policy for service egress and MC-MLPPP SAP egress. Each policy can have up to eight ingress queues and eight egress queues per service. The default policies can be copied and modified but they cannot be deleted. The default policies are identified as policy ID 1.



Note: Throughout this guide, the terms service ingress/egress and access ingress/egress are interchangeable. This chapter uses the term service ingress/egress.

The eight ingress queues can be designated as a unicast, broadcast, multicast, or unknown queue for the purposes of FC-to-queue mapping.

The default policies are applied to the appropriate interface, by default. For example, the default service ingress policy is applied to access ingress SAPs. The default service egress policy is applied to access egress SAPs and MC-MLPPP egress SAPs. You must explicitly associate other QoS policies.

6.2 Basic configuration

This section contains the following topics related to creating and applying service ingress and service egress QoS policies:

- [Creating service egress and ingress QoS policies](#)
- [Applying service egress and ingress policies](#)
- [Default service egress and ingress policy values](#)

A basic service egress QoS policy must conform to the following:

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

A basic service ingress QoS policy must conform to the following:

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

6.2.1 Creating service egress and ingress QoS policies

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP or IP interface, a default QoS policy is applied. Perform the following to configure a QoS policy:

- [Creating a service egress QoS policy](#)
- [Creating a service ingress QoS policy](#)
- [Creating an MC-MLPPP SAP egress QoS policy](#)

6.2.1.1 Creating a service egress QoS policy

After the policy is created, the policy's behavior can be defined. In addition, the behavior of policy's FC and queue can be changed from their default settings:

- [Creating a service egress QoS forwarding class](#)
- [Creating a service egress QoS queue](#)

Define the following attributes to create a service egress policy:

- a unique policy ID value – the system does not dynamically assign a value. Optionally, a policy name can be configured after the policy has been created.
- a default queue for the service egress policy
- the scope – the service egress policy must be defined as having either an **exclusive** scope for one-time use or a **template** scope that enables its use with multiple SAPs

Use the following CLI syntax to create a service egress QoS policy:

CLI syntax:

```
config>qos
  sap-egress policy-id
    description description-string
    policy-name policy-name
    queue queue-id [queue-type]
    scope {exclusive|template}
```

Example:

```
*A:ALU-1>configure qos sap-egress 600 create
config>qos>sap-egress$ fc be create
config>qos>sap-egress>fc$ exit
config>qos>sap-egress$ policy-name "sap_egr_87"
config>qos>sap-egress$ queue 2 expedite create
config>qos>sap-egress>queue$ exit
config>qos>sap-egress# scope exclusive
config>qos>sap-egress# exit
```

The following output displays the service egress policy 600 configuration:

```
*A:ALU-1 config>qos# info
```

```

-----
#-----
echo "QoS Policy Configuration"
#-----
...
    sap-egress 600 create
        scope exclusive
        policy-name "sap_egr_87"
        queue 1 create
        exit
        queue 2 expedite create
        exit
        fc be create
        exit
    exit
-----
*A:ALU-1

```

6.2.1.1.1 Creating a service egress QoS forwarding class

Define the following attributes to create a service egress forwarding class:

- the egress dot1p priority bits value
- the DSCP name and DSCP priority bits mapping

Optionally, you can enter a *queue-id* value to override the default forwarding class-to-queue mapping for the egress policy. The queue must exist before it can be associated with an FC.

Use the following CLI syntax to create a service egress forwarding class:

CLI syntax:

```

config>qos
  sap-egress policy-id
    fc fc-name
      dscp dscp-name
      dscp in-profile dscp-name out-profile dscp-name
      dot1p dot1p-value
      dot1p in-profile dot1p-value out-profile dot1p-value
      queue queue-id

```

Example:

```

*A:ALU-1>config>qos# sap-egress 600 fc be create
config>qos>sap-egress>fc# dscp cpl
config>qos>sap-egress>fc# dot1p in-profile 2 out-profile 3
config>qos>sap-egress>fc# exit
config>qos# exit
*A:ALU-1#

```

The following output displays the forwarding class configuration for service egress policy 600:

```

*A:ALU-1>config>qos# info
-----
#-----
echo "QoS Policy Configuration"
#-----
....
    sap-egress 600 create
        scope exclusive
        queue 1 create

```



```

        exit
        queue 2 expedite create
        exit
        fc be create
            dot1p in-profile 2 out-profile 3
            dscp cp1
        exit
    exit
-----
*A: ALU-1

```

6.2.1.1.2 Creating a service egress QoS queue

Define the following attributes to create a service egress queue:

- adaptation-rule – the method used by the system to derive the PIR and CIR for the queue
- cbs – overrides the reserved buffers default for the queue
- high-prio-only – the percentage of buffer space for the queue to be used exclusively by in-profile packets
- mbs – the maximum number of buffers allowed for a specific queue
- rate – the PIR and CIR values for the queue
- slope-policy – the slope policy for the queue

Use the following CLI syntax to configure the service egress QoS queue parameters:

CLI syntax:

```

config>qos
  sap-egress policy-id
    queue queue-id [queue-type]
      adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
      cbs size-in-kbytes
      high-prio-only percent
      mbs size [bytes | kilobytes]
      rate pir-rate [cir cir-rate]
      slope-policy name

```

Example:

```

*A:ALU-1# configure qos sap-egress 500
config>qos>sap-egress# queue 7
config>qos>sap-egress>queue# adaptation-rule pir closest cir closest
config>qos>sap-egress>queue# cbs 10
config>qos>sap-egress>queue# high-prio-only 10
config>qos>sap-egress>queue# mbs 10
config>qos>sap-egress>queue# rate max cir max
config>qos>sap-egress>queue# slope-policy "Slope Policy"
config>qos>sap-egress>queue# exit
config>qos>sap-egress# exit
*A:ALU-1#

```

The following output displays the queue configuration for service egress policy 500:

```

ALU-1>config>qos# info
-----
#-----
echo "QoS Policy Configuration"
#-----

```

```

....
    sap-egress 500 create
      description "Egress Policy 500"
      queue 1 create
      exit
      queue 7 best-effort create
        rate max cir max
        cbs 10
        mbs 10
        high-prio-only 10
      exit
      fc be create
      exit
      fc ef create
        dscp in-profile cp2 out-profile cp3
      exit

```

6.2.1.2 Creating a service ingress QoS policy

After the policy is created, the policy's behavior can be defined. In addition, the behavior of policy's FC and queue can be changed from their default settings:

- [Creating a service ingress forwarding class](#)
- [Creating a service ingress QoS queue](#)

To create an service ingress policy, define the following:

- a policy ID value – the system does not dynamically assign a value. Optionally, a policy name can be configured after the policy has been created.
- a description – provides a brief overview of policy features
- a default forwarding class for the policy – all packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class
- a default priority for all packets received on an ingress SAP using this policy
- the dot1p parameters – this configuration creates a mapping between the dot1p bits of the ingress traffic and the forwarding class
- the DSCP parameters – this configuration creates a mapping between the DSCP of the ingress traffic and the forwarding class
- the forwarding class parameters – overrides the default forwarding class for the policy by assigning the forwarding class to one or more of the following queue designations: broadcast-queue, multicast-queue, unknown-queue, or (unicast) queue (see [Creating a service ingress forwarding class](#))

A service 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.

Use the following CLI syntax to create a service ingress QoS policy:

CLI syntax:

```

config>qos
  sap-ingress policy-id
    description description-string
    default-fc fc-name
    default-priority {low|high}
    dot1p dot1p-priority fc fc-name priority {high|low}
    dscp dscp-name fc fc-name priority {high|low}

```

```
policy-name policy-name
```

Example:

```
*A:ALU-1>config>qos#
config>qos# sap-ingress 100 create
config>qos>sap-ingress$ description "Used on VPN SAP"
config>qos>sap-ingress$ default-fc be
config>qos>sap-ingress$ default-priority low
config>qos>sap-ingress$ dot1p 1 fc be priority low
config>qos>sap-ingress$ dscp be fc be priority low
config>qos>sap-ingress$ policy-name "sap_ingr_15"
config>qos>sap-ingress$ exit
config>qos# exit
*A:ALU-1#
```

The following output displays the configuration for service ingress policy 100:

```
ALU-1>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
....
    sap-ingress 100 create
        description "Used on VPN SAP"
        queue 1 priority-mode create
        exit
        dot1p 1 fc "be" priority low
        dscp be fc "be" priority low
        policy-name "sap_ingr_15"
    exit
...
-----
```

6.2.1.2.1 Creating a service ingress forwarding class

Use the following syntax to define a service ingress forwarding class that overrides the default forwarding type that is defined by the [default-fc](#) command. The queue must exist before it can be associated with an FC.

CLI syntax:

```
config>qos>sap-ingress policy-id
    fc fc-name
        broadcast-queue queue-id
        queue queue-id
        multicast-queue queue-id
        unknown-queue queue-id
```

Example:

```
*A:ALU-1# config>qos# sap-ingress 100 fc af create
config>qos>sap-ingress>fc# queue 2
config>qos>sap-ingress>fc# broadcast-queue 3
config>qos>sap-ingress>fc# multicast-queue 3
config>qos>sap-ingress>fc# unknown-queue 3
config>qos>sap-ingress>fc# exit
config>qos# exit
*A:ALU-1#
```

The following output displays the forwarding class override value configuration for service ingress policy 100:

```
ALU-1>config>qos# info
#-----
#-----
echo "QoS Policy Configuration"
#-----
      sap-ingress 100 create
...
      fc "af" create
        queue 2
        broadcast-queue 3
        multicast-queue 3
        unknown-queue 3
      exit
      fc "ef" create
        queue 4
        broadcast-queue 5
        multicast-queue 5
        unknown-queue 5
      exit
    exit
...
#-----
```

6.2.1.2.2 Creating a service ingress QoS queue

To create service ingress queue parameters, define the following:

- a new queue ID value – the system does not dynamically assign a value
- the queue parameters – ingress queues support explicit and auto-expedite hardware queue scheduling as well as a configurable queue mode

Use the following CLI syntax to configure SAP ingress QoS queue parameters:

CLI syntax:

```
config>qos# sap-ingress policy-id
      queue queue-id [queue-type][queue-mode]
        adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
        cbs size-in-kbytes
        high-prio-only percent
        mbs size [bytes | kilobytes]
        rate pir-rate [cir cir-rate]
        slope-policy name
```

Example:

```
*A:ALU-1# configure qos sap-ingress 100 queue 2 create
config>qos>sap-ingress>queue$ adaptation-rule pir closest   cir closest
config>qos>sap-ingress>queue$ cbs 1500
config>qos>sap-ingress>queue$ high-prio-only 10
config>qos>sap-ingress>queue$ mbs 10
config>qos>sap-ingress>queue$ rate 2500 cir 2500
config>qos>sap-ingress>queue$ slope-policy "SlopePolicyIngress"
config>qos>sap-ingress>queue$ exit
```

The following output displays the queue configuration for service ingress policy 100:

```
ALU-1>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
...
    sap-ingress 100 create
        description "Used on VPN SAP"
        queue 1 priority-mode create
        exit
        queue 2 priority-mode create
            rate 2500 cir 2500
            mbs 10
            cbs 1500
            high-prio-only 10
    ...
#-----
```

6.2.1.3 Creating an MC-MLPPP SAP egress QoS policy

After the policy is created (syntax below), the policy's FC and queue behavior can be defined:

- [Creating an MC-MLPPP SAP egress QoS forwarding class](#)
- [Creating an MC-MLPPP SAP egress QoS queue](#)

Define the following attributes to create an MC-MLPPP SAP egress policy:

- a unique policy ID value – the system does not dynamically assign a value. Optionally, a policy name can be configured after the policy has been created.
- a default queue for the MC-MLPPP SAP egress policy

Use the following CLI syntax to create an MC-MLPPP SAP egress QoS policy:

CLI syntax:

```
config>qos>mc-mlppp
    sap-egress policy-id
        description description-string
        fc fc-name
        policy-name policy-name
        queue queue-id
```

Example:

```
*A:ALU-1>configure qos>mc-mlppp# sap-egress 300 create
config>qos>mc-mlppp>sap-egress$ fc be create
config>qos>mc-mlppp>sap-egress>fc$ exit
config>qos>mc-mlppp>sap-egress>fc$ exit
config>qos>mc-mlppp>sap-egress$ policy-name "sap_egr_mc_2"
config>qos>mc-mlppp>sap-egress>queue$ exit
config>qos>mc-mlppp>sap-egress# exit
*A:ALU-1#
```

The following output displays the MC-MLPPP SAP egress policy 300 configuration:

```
*A:ALU-1 config>qos# info
#-----
echo "QoS Policy Configuration"
```

```
#-----
...
    sap-egress 300 create
        queue 1 create
        exit
        queue 2 create
        exit
        fc be create
        policy-name "sap_egr_mc_2"
        exit
    exit
```

6.2.1.3.1 Creating an MC-MLPPP SAP egress QoS forwarding class

Define the following attributes to create an MC-MLPPP SAP egress forwarding class:

- the DSCP name and DSCP priority bits mapping

Optionally, you can enter a *queue-id* value to override the default forwarding class-to-queue mapping for the egress policy. The queue must exist before it can be associated with an FC.

Use the following CLI syntax to create an MC-MLPPP SAP egress forwarding class:

CLI syntax:

```
config>qos>mc-mlppp
    sap-egress policy-id
        fc fc-name
        dscp dscp-name
        queue queue-id
```

Example:

```
*A:ALU-1>config>qos>mc-mlppp# sap-egress 300 fc be create
config>qos>mc-mlppp>sap-egress>fc# dscp af13
config>qos>mc-mlppp>sap-egress>fc# exit
config>mc-mlppp>qos# exit
*A:ALU-1#
```

The following output displays the forwarding class configuration for MC-MLPPP SAP egress policy 300:

```
*A:ALU-1>config>qos# info
-----
#-----
echo "QoS Policy Configuration"
#-----
....
    sap-egress 300 create
        queue 1 create
        exit
        queue 2 create
        exit
        fc be create
            dscp af13
        exit
    exit
-----
*A: ALU-1
```

6.2.1.3.2 Creating an MC-MLPPP SAP egress QoS queue

Define the following attributes to create an MC-MLPPP SAP egress queue:

- adaptation-rule – the method used by the system to derive the PIR for the queue
- cbs – overrides the reserved buffers default for the queue
- high-prio-only – the percentage of buffer space for the queue to be used exclusively by in-profile packets
- mbs – the maximum amount of buffer space allowed for a specific queue
- rate – the PIR value for the queue
- slope-policy – the slope policy for the queue

Use the following CLI syntax to configure the MC-MLPPP SAP egress QoS queue parameters:

CLI syntax:

```
config>qos>mc-mlppp
  sap-egress policy-id
    queue queue-id
      adaptation-rule [pir adaptation-rule]
      cbs size-in-kbytes
      high-prio-only percent
      mbs size [bytes | kilobytes]
      rate pir-rate
      slope-policy name
```

Example:

```
*A:ALU-1# configure qos mc-mlppp sap-egress 300
config>qos>mc-mlppp>sap-egress# queue 7
config>qos>mc-mlppp>sap-egress>queue# adaptation-rule pir closest
config>qos>mc-mlppp>sap-egress>queue# cbs 10
config>qos>mc-mlppp>sap-egress>queue# high-prio-only 10
config>qos>mc-mlppp>sap-egress>queue# mbs 10
config>qos>mc-mlppp>sap-egress>queue# rate max
config>qos>mc-mlppp>sap-egress>queue# slope-policy "Slope Policy"
config>qos>mc-mlppp>sap-egress>queue# exit
config>qos>mc-mlppp>sap-egress# exit
*A:ALU-1#
```

The following output displays the queue configuration for MC-MLPPP SAP egress policy 300:

```
ALU-1>config>qos# info
-----
#-----
echo "QoS Policy Configuration"
#-----
....
    sap-egress 300 create
      description "Egress Policy 300"
      queue 1 create
      exit
      queue 7 best-effort create
        rate max
        cbs 10
        mbs 10
        high-prio-only 10
      exit
```

```

        fc be create
        exit
        fc ef create
            dscp af13
        exit
    exit
...
-----
ALU-1#

```

6.2.2 Applying service egress and ingress policies

Apply service egress and ingress policies to the following service SAPs:

- VLL services
 - Epipe
 - Cpipe
 - Apipe
 - Apipe SAPs that are members of a VCC SAP aggregation group
 - Fpipe
 - Hpipe
 - Lpipe
- VPLS
- IES
- VPRN

See the 7705 SAR Services Guide, "Service Overview", for information about configuring service parameters.

6.2.2.1 VLL and VPLS services

Applying QoS policies is done in the same way for service ingress and service egress VPLS and VLL SAPs. The following example shows how to apply QoS policies to Epipe SAPs. Use the policy ID number or name to identify the policy.

CLI syntax:

```

config>service>epipe service-id customer customer-id
    sap sap-id
        egress
            qos policy-id
        ingress
            qos policy-id

```

The following output displays an Epipe service configuration with service ingress policy 100 and service egress policy 105 applied to the SAP.

```

ALU-1>config>service# info
#-----
echo "QoS Policy Configuration"
#-----

```



```

...
    epipe 6 customer 6 vpn 6 create
        description "Distributed Epipe service to west coast"
        sap 1/1/10:0 create
            ingress
                qos 100
            exit
            egress
                qos 105
            exit
        exit
    spoke-sdp 2:6 create
        ingress
            vc-label 6298
        exit
        egress
            vc-label 6300
        exit
    exit
    no shutdown
exit
...
#-----
ALU-1>config>service#

```

6.2.2.2 IES and VPRN services

Applying QoS policies is done in the same way for service ingress and service egress IES and VPRN SAPs. The following example shows how to apply QoS policies to IES SAPs. Use the policy ID number or name to identify the policy.

CLI syntax:

```

config>service>ies service-id customer customer-id
    interface ip-interface-name
        sap sap-id
            egress
                qos policy-id
            ingress
                qos policy-id

```

The output examples for IES and VPRN services are similar to the example shown previously for an Epipe service configuration.

6.2.3 Default service egress and ingress policy values

The default service egress and ingress policies are identified as policy-id 1 and have the default policy name "default". The default policies cannot be edited or deleted. The following sections display default policy parameters:

- [Service egress policy defaults](#)
- [Service ingress policy defaults](#)

6.2.3.1 Service egress policy defaults

The following output shows the service egress policy defaults.

```
ALU-1>config>qos>info detail
#-----
#-----
echo "QoS Policy Configuration"
#-----
...
    sap-egress 1 create
        policy-name "default"
        description "Default SAP egress QoS policy."
        scope template
        queue 1 auto-expedite create
            adaptation-rule pir closest cir closest
            rate max cir 0
            cbs default
            mbs default
            high-prio-only default
            slope-policy "default"
        ...
#-----
```

The following table lists the service egress policy defaults.

Table 43: Service egress policy defaults

Field	Default
policy-name	"default"
description	Default SAP egress QoS policy
scope	template
queue	id = 1, type = auto-expedite
adaptation-rule	pir = closest, cir = closest
rate	pir = max, cir = 0
cbs	default (8 kB for 512 byte buffer size, 18 kB for 2304 byte buffer size)
mbs	default (180 kB)
high-prio-only	default (10%)
slope-policy	default

6.2.3.2 Service ingress policy defaults

The following output shows the service ingress policy defaults.

```
ALU-1>config>qos>info detail
#-----
#-----
echo "QoS Policy Configuration"
#-----
...
    sap-ingress 1 create
        policy-name "default"
        description "Default SAP ingress QoS policy."
        scope template
        queue 1 priority-mode auto-expedite create
            adaptation-rule pir closest cir closest
            rate max cir 0
            mbs default
            cbs default
            high-prio-only default
            slope-policy "default"
        exit
        default-fc "be"
        default-priority low
    exit
...
#-----
```

The following table lists the service ingress policy defaults.

Table 44: Service ingress policy defaults

Field	Default
policy-name	"default"
description	Default SAP ingress QoS policy
scope	template
queue	id = 1, mode = priority-mode, type = auto-expedite
adaptation-rule	pir = closest, cir = closest
rate	pir = max, cir = 0
cbs	default (8 kB for 512 byte buffer size, 18 kB for 2304 byte buffer size)
mbs	default (180 kB)
high-prio-only	default (10%)
slope-policy	default
default-fc	be

Field	Default
default-priority	low

6.3 Service management tasks

This section describes the following service management tasks:

- [Deleting QoS policies](#)
- [Copying and overwriting QoS policies](#)
- [Editing QoS policies](#)

6.3.1 Deleting QoS policies

Every service SAP is associated, by default, with the appropriate service egress or ingress policy (policy-id 1). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the SAP configuration. When you remove a non-default service egress or ingress policy, the association reverts to the default policy-id 1.

A QoS policy cannot be deleted until it is removed from all SAPs where it is applied.

6.3.1.1 Removing a QoS policy from a service SAP

Use the following syntax to remove a QoS policy from an Epipe service SAP. The syntax for Apipe, Cpipe, Fpipe, Hpipe, and Lpipe service SAPs is similar.

CLI syntax:

```
config>service>epipe service-id customer customer-id
  sap sap-id
    egress
      no qos policy-id
    ingress
      no qos policy-id
```

Example:

```
config>service>epipe# sap 1/1/10:0
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# no qos
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# egress
config>service>epipe>sap>egress# no qos
config>service>epipe>sap>egress# exit
```

The following Epipe service output shows that the SAP service egress and ingress reverted to policy-id "1" when the non-default policies were removed from the configuration.

```
ALU-1>config>service>epipe# info detail
-----
description "Distributed Epipe service to west coast"
service-mtu 1514
```

```

    sap 1/1/10:0 create
    no description
    no multi-service-site
    ingress
        qos 1
        exit
    egress
        qos 1
        exit
    no collect-stats
    no accounting-policy
    no shutdown
exit
spoke-sdp 2:6 vc-type ether create
    no shutdown
exit
no shutdown
-----
ALU-1>config>service>epipe#

```

6.3.1.2 Removing a policy from the QoS configuration

Use the following syntax to remove a QoS policy:

CLI syntax:

```
config>qos# no sap-ingress policy-id
```

6.3.2 Copying and overwriting QoS policies

You can copy an existing service egress or ingress policy, rename it with a new policy ID value, or overwrite an existing policy ID. The **overwrite** option must be specified or an error occurs if the destination policy ID exists.

Use the following syntax to overwrite an existing QoS policy ID:

CLI syntax:

```
config>qos# copy sap-ingress source-policy-id dest-policy-id overwrite
```

Example:

```

config>qos# copy sap-ingress 100 200
config>qos# copy sap-ingress 200 101
MINOR: CLI Destination "101" exists use {overwrite}
config>qos# copy sap-ingress 200 101 overwrite
config>qos#

```

The following output displays the copied policies:

```

ALU-1>config>qos# info
-----
...
exit
    sap-ingress 100 create
    description "Used on VPN sap"
    queue 1 create
    exit
    rate 11000

```

```
...      exit
...      sap-ingress 101 create
...      description "Used on VPN sap"
...      queue 1 create
...      exit
...      rate 11000
...      exit
...      sap-ingress 200 create
...      description "Used on VPN sap"
...      queue 1 create
...      exit
...      rate 11000
...      exit
...
-----
ALU-1>config>qos#
```

6.3.3 Editing QoS policies

You can change existing QoS policies and entries in the CLI. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors, copy the policy to a work area, make the edits, and then write over the original policy.

6.4 Service egress and ingress QoS policy command reference

6.4.1 Command hierarchies

- [Service egress QoS policy configuration commands](#)
- [Service ingress QoS policy configuration commands](#)
- [MC-MLPPP SAP egress QoS policies](#)
- [Operational commands](#)
- [Show commands](#)

6.4.1.1 Service egress QoS policy configuration commands

```

config
- qos
- [no] sap-egress policy-id [create]
  - description description-string
  - no description
  - [no] fc fc-name [create]
    - dot1p {dot1p-value | in-profile dot1p-value out-profile dot1p-value}
    - no dot1p
    - dscp dscp-name
    - dscp in-profile dscp-name out-profile dscp-name
    - no dscp
    - queue queue-id
    - no queue
  - packet-byte-offset [add bytes | subtract bytes | none]
  - no packet-byte-offset
  - policy-name policy-name
  - no policy-name
  - queue queue-id [queue-type] [create]
  - no queue queue-id
    - adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
    - no adaptation-rule
    - cbs {size-in-kbytes | default}
    - no cbs
    - high-prio-only percent
    - no high-prio-only
    - mbs size [bytes | kilobytes]
    - no mbs
    - packet-byte-offset [add bytes | subtract bytes | none]
    - no packet-byte-offset
    - rate pir-rate [cir cir-rate]
    - no rate
    - slope-policy name
    - no slope-policy
  - scope {exclusive | template}
  - no scope

```

6.4.1.2 Service ingress QoS policy configuration commands

```

config

```

```

- qos
- [no] sap-ingress policy-id [create]
- default-fc fc-name
- no default-fc
- default-priority {high | low}
- no default-priority
- description description-string
- no description
- dot1p dot1p-priority [fc fc-name] [priority {high | low}]
- no dot1p dot1p-priority
- dscp dscp-name [dscp-name...(up to 8 max)] [fc fc-name] [priority {high | low}]
- no dscp dscp-name [dscp-name...(up to 8 max)]
- [no] fc fc-name [create]
- broadcast-queue queue-id
- no broadcast-queue
- de-l-out-profile
- no de-l-out-profile
- multicast-queue queue-id
- no multicast-queue
- profile {in | out}
- no profile
- queue queue-id
- no queue
- unknown-queue queue-id
- no unknown-queue
- packet-byte-offset [add bytes | subtract bytes | none]
- no packet-byte-offset
- policy-name policy-name
- no policy-name
- queue queue-id [queue-type] [queue-mode] [create]
- no queue queue-id
- adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
- no adaptation-rule
- cbs {size-in-kbytes | default}
- no cbs
- high-prio-only percent
- no high-prio-only
- mbs size [bytes | kilobytes]
- no mbs
- packet-byte-offset [add bytes | subtract bytes | none]
- no packet-byte-offset
- rate pir-rate [cir cir-rate]
- no rate
- slope-policy name
- no slope-policy
- scope {exclusive | template}
- no scope

```

6.4.1.3 MC-MLPPP SAP egress QoS policies

```

config
- qos
- mc-mlppp
- [no] sap-egress policy-id [create]
- description description-string
- no description
- [no] fc fc-name [create]
- dscp dscp-name
- no dscp
- queue queue-id
- no queue

```



```

- policy-name policy-name
- no policy-name
- [no] queue queue-id [create]
  - adaptation-rule pir adaptation-rule
  - no adaptation-rule
  - cbs {size-in-kbytes | default}
  - no cbs
  - high-prio-only percent
  - no high-prio-only
  - mbs size [bytes | kilobytes]
  - no mbs
  - rate pir-rate
  - no rate
  - slope-policy name
  - no slope-policy
- scope {exclusive | template}
- no scope

```

6.4.1.4 Operational commands

```

config
- qos
  - copy sap-egress src-pol dst-pol [overwrite]
  - copy sap-ingress src-pol dst-pol [overwrite]

```

6.4.1.5 Show commands

```

show
- qos
  - sap-egress [policy-id] [standard | mc-mlppp] [association | match-criteria | detail]
  - sap-egress summary
  - sap-ingress [policy-id] [association | match-criteria | detail]
  - sap-ingress summary
- pools mda mda-id [detail]
- pools mda mda-id [egress | ingress] [ring | v-port] [detail]

```

6.4.2 Command descriptions

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

6.4.2.1 Configuration commands

- [Generic commands](#)
- [Service egress QoS policy commands](#)
- [Service egress QoS policy forwarding class commands](#)
- [MC-MLPPP SAP egress QoS policy commands](#)
- [MC-MLPPP forwarding class commands](#)
- [MC-MLPPP queue commands](#)
- [Service ingress QoS policy commands](#)
- [Service ingress QoS policy forwarding class commands](#)
- [Service queue QoS policy commands](#)

6.4.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>sap-egress

config>qos>sap-ingress

config>qos>mc-mlppp>sap-egress

Description

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

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

Default

n/a

Parameters

description-string

a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

6.4.2.1.2 Service egress QoS policy commands

sap-egress

Syntax

[no] **sap-egress** *policy-id* [create]

Context

config>qos

Description

This command is used to create or edit a service egress QoS policy. The egress policy defines the service-level agreement (SLA) for service packets as they egress on the SAP.

Policies in effect are templates that can be applied to multiple services as long as the **scope** of the policy is **template**. The queues defined in the policy are not instantiated until a policy is applied to a service.

A sap-egress policy differs from sap-ingress policies in the complexity of the QoS parameters that can be defined. At ingress, policies determine queue mappings based on ingress DSCP and dot1p criteria. Multiple queues can be created per forwarding class and each queue can have different CIR or PIR parameters.

At egress, the policies are much simpler, since the forwarding class and in-profile or out-of-profile determination is done at the original service ingress SAP. Egress SAP QoS policies allow the definition of queues and the mapping of forwarding classes to those queues. Each queue needs to have a relative CIR for determining its allocation of QoS resources during periods of congestion. A PIR can also be defined that forces a hard limit on the packets transmitted through the queue. When the forwarding class is mapped to the queue, a dot1p value can optionally be specified. If specified, and the SAP has a dot1q or qinq encapsulation type, the dot1p value will be used for all packets that egress on that forwarding class. If the SAP is qinq-encapsulated, the **qinq-mark-top-only** command (under **config>service**) can be used to specify which qtags will have their dot1p marked or re-marked with the specified dot1p value. If the dot1p value is not specified, a dot1p value of 0 will be used. If the SAP is null-encapsulated, the dot1p value has no meaning.



Note: The ATM access egress shaping configuration in a SAP egress QoS policy is ignored when that policy is assigned to an ATM SAP. The shaping of the egress cell stream is controlled by the **atm-td-profile** command. If the **atm-td-profile** is not configured, the default **atm-td-profile** is in effect. See the [atm-td-profile](#) command for more information.

The sap-egress policy with *policy-id* 1 is the default sap-egress QoS policy and is applied to service egress SAPs when an explicit policy is not specified or removed.

The factory default settings for **sap-egress** *policy-id* 1 define a single queue with PIR set to the maximum value and a CIR set to 0. The single queue is the default queue and all forwarding classes will map to it. Packets being tagged according to the defined SAP encapsulation will have the dot1p bits set to 0 for the new tags being added. If the tag already exists and the default **sap-egress** policy is being used, the dot1p bits are not changed.



Note: If the egress port encapsulation type is qinq, the SAP type is X.Y, and one tag already exists and a new tag must be added, then the new outer tag's dot1p bits will be set to the inner tag's dot1p bits value. For all other port types and SAP types, new tags will have a dot1p-bits value of 0 if the default policy is used.

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

The **no** form of this command deletes the sap-egress policy. A policy cannot be deleted until it is removed from all service SAPs where it is applied. When a sap-egress policy is removed from a SAP, the SAP will revert to the default **sap-egress** *policy-id* 1.

Parameters

policy-id

uniquely identifies the policy. The *policy-name* cannot be used to create the policy, but it can be used to reference a policy that already exists.

Values 1 to 65535, or *policy-name* (up to 64 characters)

Default n/a

create

keyword used to create a service egress QoS policy

fc

Syntax

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

Context

config>qos>sap-egress

Description

The **fc** *fc-name* mode within the SAP egress QoS policy is used to contain the explicitly defined queue mapping and dot1p marking commands for *fc-name*. When the mapping for *fc-name* points to the default queue and the dot1p marking is not defined, the mode for *fc-name* is not displayed in the **show configuration** or **save configuration** output unless the detail option is specified.

The **no** form of the command removes the explicit queue mapping and dot1p marking commands for *fc-name*. The queue mapping reverts to the default queue for *fc-name*, and the dot1p marking (if appropriate) uses the default of 0.

Default

n/a

Parameters

fc-name

specifies the forwarding class queue mapping or dot1p marking is to be edited. The value given for *fc-name* must be one of the predefined forwarding classes in the system.

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

create

keyword used to create a SAP egress forwarding class policy

packet-byte-offset

Syntax

packet-byte-offset [**add bytes** | **subtract bytes** | **none**]

no packet-byte-offset

Context

config>qos>sap-egress

config>qos>sap-egress>queue

config>qos>sap-ingress

config>qos>sap-ingress>queue

Description

This command is used to modify the size of the packet that schedulers operate on. Modification only impacts schedulers and queue statistics. The actual packet size is not modified, nor can it be. Only the size used by the schedulers to determine the scheduling is changed. The **packet-byte-offset** command is meant to be a mechanism that can be used to compensate for downstream encapsulation or header removal. The scheduling rates are affected by the offset, as well as the statistics (accounting) associated with the queue. The **packet-byte-offset** command does not affect port-level and service-level statistics. It only affects the queue statistics. The network-queue policy applies in both the ingress and egress directions.

The **add** and **subtract** keywords are mutually exclusive. Either **add**, **subtract**, or **none** must be specified.

There are three possible modes of **packet-byte-offset** operation:

- **no packet-byte-offset** – enables legacy behavior so that no modification is performed
- **packet-byte-offset** – automatic adjustment mode. Rates apply to packets based on the received packet size at ingress (this is also known as packet size on the wire, less the Layer 1 headers, the inter-frame GAP and the Preamble) and to the transmitted packet size at egress, which includes 4 bytes of

Ethernet FCS. At ingress, all internal headers and associated service headers are discounted during scheduling operation. At egress, 4 bytes are added to accommodate for Ethernet FCS.

- **packet-byte-offset** [add *bytes* | subtract *bytes*] – automatic correction followed by addition or subtraction of a specified number of bytes. This command first performs the **packet-byte-offset** operation as captured above and then adds or subtracts a specific number of bytes. Rates apply to packets based on the size of the packet at the ingress or egress port plus or minus an offset.

Packet byte offset configuration can be applied at the policy level, in which case it applies to all of the queues within the policy, or at the individual queue level so that it applies only to a specific queue.

The **no** version of this command enables legacy 7705 SAR behavior where the queue rates are relative to the packet size with the internal fabric header added, but without the FCS.

Parameters

add bytes

after automatic adjustment for internal headers (for example, added FCS or removal of internal service/overhead), adds the specified number of bytes to each packet associated with the queue for scheduling and accounting purposes. From the queue's perspective, the packet size is increased by the amount being added to each packet.

Values 2 to 62, in steps of 2

subtract bytes

after automatic adjustment for internal headers (for example, added FCS or removal of internal service/overhead), subtracts the specified number of bytes from each packet associated with the queue for scheduling and accounting purposes. From the queue's perspective, the packet size is reduced by the amount being subtracted from each packet.

Values 2 to 62, in steps of 2

none

the packet size is left unchanged

policy-name

Syntax

policy-name *policy-name*

no *policy-name*

Context

config>qos>sap-egress

Description

This command configures a policy name for the SAP egress policy.

Parameters

policy-name

specifies the name for the policy, up to 64 characters

queue

Syntax

```
queue queue-id [queue-type] [create]
no queue queue-id
```

Context

```
config>qos>sap-egress
```

Description

This command enables the context to configure a service egress policy queue. Explicit definition of an egress queue’s hardware scheduler status is supported. A single egress queue allows support for multiple forwarding classes.

The default behavior automatically chooses the expedited or non-expedited nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (**nc**, **ef**, **h1**, or **h2**), the queue is treated as an expedited queue by the hardware schedulers. When any non-expedited forwarding classes are mapped to the queue (**be**, **af**, **l1**, or **l2**), the queue is treated as best effort (**be**) by the hardware schedulers. The expedited hardware schedulers are used to enforce expedited access to egress ports. The hardware status of the queue must be defined at the time of queue creation within the policy.

The no form of the command removes the *queue-id* from the service egress policy. Removing the *queue-id* also removes it from any existing SAPs using the policy. If any forwarding classes are mapped to the queue, they revert to the default queue.

When a queue is removed, pending accounting information for each service egress queue created due to the definition of the queue in the policy is discarded.

Default

```
n/a
```

Parameters

queue-id

the *queue-id* for the service egress queue, expressed as a decimal integer. The *queue-id* uniquely identifies the queue within the policy.

Values 1 to 8

Default n/a

queue-type

the **expedite**, **best-effort** and **auto-expedite** queue types are mutually exclusive. Each defines the method that the system uses to service the queue from a hardware perspective. A keyword must be specified at the time the queue is created in the service egress policy. If an attempt is made to change the keyword after the queue is initially defined, an error is generated.

expedite

the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue

best-effort

the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue

auto-expedite

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

Values expedite, best-effort, auto-expedite

Default auto-expedite

create

keyword used to create a service egress queue policy

scope**Syntax**

scope {**exclusive** | **template**}

no scope

Context

config>qos>sap-egress

Description

This command is used to enter the scope of the policy. The scope of the policy cannot be changed if the policy is applied to one or more services.

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

Default

template

Parameters**exclusive**

when the scope of a policy is defined as **exclusive**, the policy can only be applied to a single SAP. Attempting to assign the policy to a second SAP will result in an error message. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.

template

when the scope of a policy is defined as **template**, the policy can be applied to multiple SAPs on the router

6.4.2.1.3 Service egress QoS policy forwarding class commands

dot1p

Syntax

```
dot1p {dot1p-value | in-profile dot1p-value out-profile dot1p-value}
no dot1p
```

Context

```
config>qos>sap-egress>fc
```

Description

This command explicitly defines the egress dot1p priority bits values for the forwarding class.



Note:

- When the **dot1p dot1p-value** command is used, the value is applied to both in-profile and out-of-profile packets. When the **dot1p in-profile dot1p-value out-profile dot1p-value** form is used, different dot1p values for in-profile or out-of-profile packets can be specified.
- For QinQ applications, the top and bottom dot1p bits or top-only dot1p bits are set based on the **qinq-mark-top-only** configuration. For more information, see the command description and the "QinQ Top-Only Mark Option" section in the 7705 SAR Services Guide.

The **no** form of the command sets the dot1p priority bits value to 0.

Default

0

Parameters

dot1p-value

the explicit dot1p value for the specified forwarding class. Setting the value to 0 is equivalent to removing the marking value.

Values 0 to 7

Values n/a

in-profile *dot1p-value*

the dot1p in-profile mapping for the specified forwarding class. Setting the value to 0 is equivalent to removing the mapping value.

Values 0 to 7

Default n/a

out-profile *dot1p-value*
the dot1p out-profile mapping for the specified forwarding class. Setting the value to 0 is equivalent to removing the mapping value.

Values 0 to 7

Default n/a

dscp

Syntax


dscp *dscp-name*
dscp in-profile *dscp-name* **out-profile** *dscp-name*
no dscp

Context

config>qos>sap-egress>fc

Description

This command defines the DSCP name or DSCP priority bits mapping for the forwarding class.

 **Note:** When the **dscp** *dscp-name* command is used, the *dscp-name* is applied to all packets regardless of the profile state. The **in-profile** and **out-profile** form of the command allows differentiated values to be applied to packets based on the profile state.

Access IP traffic (that is, VPRN and IES access interfaces) will always be re-marked if the interfaces are configured as part of the SAP egress QoS policy.

The **no** form of the command removes the DSCP mapping associated with the forwarding class.

Default

n/a

Parameters

dscp-name
a system-defined, case-sensitive name

Values A maximum of 64 DSCP rules are allowed on a single policy. The specified name must exist as a *dscp-name*. The following table lists all the valid DSCP names.

Table 45: Valid DSCP names

dscp-name
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

in-profile *dscp-name*
the DSCP in-profile mapping for the specified forwarding class

Values any name listed in [Table 45: Valid DSCP names](#)

out-profile *dscp-name*
the DSCP out-profile mapping for the specified forwarding class

Values any name listed in [Table 45: Valid DSCP names](#)

queue

Syntax

queue *queue-id*

no queue

Context

config>qos>sap-egress>fc

Description

This command specifies the egress queue to which the traffic associated with the forwarding class is to be forwarded. The command overrides the default queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy before the mapping can be made. When the forwarding class mapping is executed, all traffic classified to the *fc-name* on a SAP using this policy will use the indicated queue.

The **no** form of the command sets the *queue-id* back to the default queue for the forwarding class (queue 1).

Default

no queue

Parameters

queue-id

the service egress *queue-id* to be associated with the forwarding class. The *queue-id* must be an existing queue defined in **sap-egress** *policy-id*.

Values 1 to 8

Default 1

6.4.2.1.4 MC-MLPPP SAP egress QoS policy commands

mc-mlppp

Syntax

mc-mlppp

Context

config>qos

Description

This command enables the context to configure MC-MLPPP SAP egress QoS commands.

sap-egress

Syntax

[no] sap-egress *policy-id* [create]

Context

config>qos>mc-mlppp

Description

This command is used to create or edit an MC-MLPPP SAP egress QoS policy. The egress policy defines the SLA for service packets as they egress on the SAP.

Policies are templates that can be applied to multiple services as long as the scope of the policy is template. The queues defined in the policy are not instantiated until a policy is applied to a service. At egress, the forwarding class and in- or out-of-profile determination is done at the original service ingress SAP. MC-MLPPP egress SAP QoS policies allow the definition of queues and the mapping of forwarding classes to those queues. Each queue must have a PIR defined that forces a hard limit on the packets transmitted through the queue.

The sap-egress policy with *policy-id* 1 is the default sap-egress QoS policy and is applied to MC-MLPPP egress SAPs when an explicit policy is not specified or is removed.

The default settings for **sap-egress *policy-id* 1** define a single queue with PIR set to the maximum value. The single queue is the default queue and all forwarding classes will map to it. Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all egress SAPs where this policy is applied. For this reason, when many changes are required on a policy, it is highly recommended that the policy be copied to a work area *policy-id*. That work-in-progress policy can be modified until

complete and then written over the original *policy-id*. Use the **config qos copy** command to maintain policies in this manner.

The **no** form of this command deletes the sap-egress policy. A policy cannot be deleted until it is removed from all service SAPs where it is applied. When a sap-egress policy is removed from a SAP, the SAP will revert to the default **sap-egress policy-id** 1, which cannot be deleted.

Parameters

policy-id

uniquely identifies the policy. The *policy-name* cannot be used to create the policy, but it can be used to reference a policy that already exists.

Values 1 to 65535, or *policy-name* (up to 64 characters)

Default n/a

create

keyword used to create an MC-MLPPP SAP egress QoS policy

fc

Syntax

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

Context

config>qos>mc-mlppp>sap-egress

Description

The **fc** *fc-name* mode within the MC-MLPPP SAP egress QoS policy is used to contain the explicitly defined queue mapping for *fc-name*.

The **no** form of the command removes the explicit queue mapping for *fc-name*. The queue mapping reverts to the default queue for *fc-name*.

Default

n/a

Parameters

fc-name

specifies that the forwarding class queue mapping is to be edited. The value given for *fc-name* must be one of the predefined forwarding classes in the system.

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

create

keyword used to create the context to configure an MC-MLPPP SAP egress forwarding class mapping queue

policy-name

Syntax

policy-name *policy-name*

no policy-name

Context

config>qos>mc-mlppp>sap-egress

Description

This command configures a policy name for the MC-MLPPP SAP egress policy.

Parameters

policy-name

specifies the name for the policy, up to 64 characters

queue

Syntax

[no] queue *queue-id* [**create**]

Context

config>qos>mc-mlppp>sap-egress

Description

This command enables the context to configure an MC-MLPPP SAP egress policy queue. Explicit definition of an egress queue's hardware scheduler status is supported. A single egress queue allows support for multiple forwarding classes.

The no form of the command removes the *queue-id* from the MC-MLPPP SAP egress policy. Removing the *queue-id* also removes it from any existing SAPs using the policy. If any forwarding classes are mapped to the queue, they revert to the default queue.

When a queue is removed, pending accounting information for each SAP egress queue created due to the definition of the queue in the policy is discarded.

Default

n/a

Parameters

queue-id

the *queue-id* for the MC-MLPPP SAP egress queue, expressed as a decimal integer. The *queue-id* uniquely identifies the queue within the policy.

Values 1 to 8

create

keyword used to create the context to configure an MC-MLPPP SAP egress policy queue

scope

Syntax

scope {**exclusive** | **template**}

no scope

Context

config>qos>mc-mlppp>sap-egress

Description

This command is used to enter the scope of the policy. The scope of the policy cannot be changed if the policy is applied to one or more services.

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

Default

template

Parameters

exclusive

when the scope of a policy is defined as **exclusive**, the policy can only be applied to a single SAP. Attempting to assign the policy to a second SAP will result in an error message. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.

template

when the scope of a policy is defined as **template**, the policy can be applied to multiple SAPs

6.4.2.1.5 MC-MLPPP forwarding class commands

dscp

Syntax

dscp *dscp-name*

no dscp

Context

config>qos>mc-mlppp>sap-egress>fc

Description

This command defines the DSCP name for the forwarding class.

Default

n/a

Parameters

dscp-name
a system-defined, case-sensitive name
Values any name listed in [Table 45: Valid DSCP names](#)

queue

Syntax

queue *queue-id*
no queue

Context

config>qos>mc-mlppp>sap-egress>fc

Description

This command specifies the MC-MLPPP egress queue to which the traffic associated with the forwarding class is to be forwarded. The command overrides the default queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy before the mapping can be made. Once the forwarding class mapping is executed, all traffic classified to the *fc-name* on a SAP using this policy will use the indicated queue. The **no** form of the command sets the *queue-id* back to the default queue for the forwarding class (queue 1).

Default

queue 1

Parameters

queue-id
the MC-MLPPP SAP egress *queue-id* to be associated with the forwarding class. The *queue-id* must be an existing queue defined in [sap-egress](#) *policy-id*.
Values 1 to 8
Default 1

6.4.2.1.6 MC-MLPPP queue commands

adaptation-rule

Syntax

adaptation-rule *pir adaptation-rule*

no adaptation-rule

Context

config>qos>mc-mlppp>sap-egress>queue

Description

This command is used to define how an operational rate is selected based on the configured PIR rate. Operational rates are the finite set of rates at which the schedulers on the network processor can operate.

The **no** form of the command removes any **adaptation-rule** constraints used to derive the operational rates for the policy. When a specific adaptation-rule is removed, the default constraints for [rate](#) apply.

Default

closest

Parameters

pir

defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* [rate](#) command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the [rate](#) command is not specified, the default applies.

adaptation-rule

specifies the constraints to be used while computing the operational PIR rate. The **max** (maximum), **min** (minimum), and **closest** parameters are mutually exclusive.

- Values**
- max** – causes the network processor to be programmed at an operational rate that is less than the configured PIR rate by up to 1.0%
 - min** – causes the network processor to be programmed at an operational rate that is greater than the configured PIR rate by up to 1.0%
 - closest** – causes the network processor to be programmed at an operational rate that is closest to the configured PIR rate

cbs

Syntax

cbs {*size-in-kbytes* | **default**}

no cbs

Context

config>qos>mc-mlppp>sap-egress>queue

Description

This command overrides the default committed buffer space (CBS) reserved buffers for the queue.

The value in kilobytes is converted automatically to the number of buffers. The conversion calculation uses a non-configurable buffer size of 2304 bytes or 512 bytes, depending on the type of adapter card. See [Table 4: Buffer support on adapter cards and platforms](#) for a list of adapter cards and their associated buffers. The calculation is:

Number of buffers = configured CBS value in bytes / buffer size in bytes

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

Default

"default" (8 kB for adapter cards and platforms with 512 byte buffer size) (18 kB for adapter cards and platforms with 2304 byte buffer size)

Parameters

size-in-kbytes

this parameter is an integer expression of the number of kilobytes reserved for the queue. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 to 131072

default

returns the CBS size to the default value

high-prio-only

Syntax

high-prio-only *percent*

no high-prio-only

Context

config>qos>mc-mlppp>sap-egress>queue

Description

The **high-prio-only** command configures the percentage of buffer space for the queue, used exclusively by high-priority packets. The specified value overrides the default value for the context.

The priority of a packet can only be set in the service ingress policy and is only applicable on the ingress queues for a SAP. The profile state is used for enqueueing priority at sap-egress.

The **no** form of this command restores the default high-priority reserved size.

Default

10 (percent)

Parameters

percent

the percentage reserved for high priority traffic on the queue

Values 0 to 100 | default 1

mbs

Syntax

mbs *size* [bytes | kilobytes]

no mbs

Context

config>qos>mc-mlppp>sap-egress>queue

Description

This command sets the Maximum Burst Size (MBS) value for the buffers of the specific queue. The value is configured in bytes or kilobytes and overrides the default value for the context. The default configuration is in kilobytes.

The **config>qos>mc-mlppp>sap-egress>info detail** screen shows the MBS in terms of bytes, unless it is a multiple of 1000. In that case, the display shows the MBS in kilobytes. For example, entering **mbs 200** or **mbs 200 kilobytes** configures and displays "200 kilobytes", entering **mbs 200000 bytes** also configures and displays "200 kilobytes", and entering **mbs 200100 bytes** configures and displays "200100 bytes".



Note: For the 7705 SAR, 1 kB of buffer management space is 1000 bytes.

The MBS value in bytes is converted automatically to packets. The conversion calculation uses a non-user-configurable buffer size of 2304 bytes or 512 bytes, depending on the type of adapter card. See [Table 4: Buffer support on adapter cards and platforms](#) for a list of adapter cards and their associated buffers. The calculation is:

Number of buffers = Configured MBS value in bytes / Buffer size in bytes (2304 or 512)

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an adapter card can exceed the total amount of buffering available. Therefore, for a packet arriving at a queue that has not exceeded its MBS size, it is not guaranteed that a buffer will be available. If a buffer is not available, the packet will be discarded. RED/WRED slope parameters can be configured to control congestion in the case where the buffer capacity of the card is becoming exhausted.

Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard against queue starvation (that is, when a queue does not receive its fair share of buffers). Another safeguard is to properly set the RED/WRED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS size assigned to the queue to the default value.

Default

180 (kB) (converted to 78 packets when buffer size is 2304 bytes and to 351 packets when buffer size is 512 bytes)

Parameters

size

the *size* parameter is an integer expression of the maximum number of bytes of buffering allowed for the queue. A value of 4000 bytes or less causes the queue to discard all packets. Selecting **default** returns the MBS to the default value.

Values 0 to 131072000 | **default**

bytes

specifies that the size entered is in bytes

kilobytes

specifies that the size entered is in kB

rate

Syntax

rate *pir-rate*

no rate

Context

config>qos>mc-mlppp>sap-egress>queue

Description

This command defines the administrative PIR parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress interface. 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 **rate** command can be executed at any time, altering the PIR rates for all queues created through the association of the MC-MLPPP SAP egress policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR parameters (max).

Default

max (this value specifies the amount of bandwidth in kb/s. The max value and the *pir-rate* value are mutually exclusive.)

Parameters

pir-rate

defines the administrative PIR rate, in kb/s, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of max is assumed. Fractional values are not allowed and must be given as a positive integer. The PIR rate has a minimum value of 8kb/s.

The actual PIR rate is dependent on the queue's [adaptation-rule](#) parameters and the actual hardware where the queue is provisioned.

Values 1 to 100000000 | max



Note: If a PIR rate lower than 8 kb/s is specified, it is rounded up to this minimum value.

slope-policy

Syntax

slope-policy *name*

no slope-policy

Context

config>qos>mc-mlppp>sap-egress queue

Description

This command specifies the slope parameters controlling the queue.

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

Default

default

Parameters

name

the name of the slope policy

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

6.4.2.1.7 Service ingress QoS policy commands

sap-ingress

Syntax

[no] sap-ingress *policy-id* [create]

Context

config>qos

Description

This command is used to create or edit the ingress policy. The ingress policy defines the SLA enforcement that service packets receive as they ingress a SAP. SLA enforcement is accomplished through the definition of queues that have FC, CIR, PIR, and MBS characteristics. The simplest policy defines a single queue that all ingress traffic flows through. Complex policies have multiple queues that indicate which queue a packet will flow through.

Policies are templates that can be applied to multiple services as long as the [scope](#) of the policy is **template**. Queues defined in the policy are not instantiated until a policy is applied to a service SAP.

It is possible that a service ingress policy will include the [dscp](#) map command and the [dot1p](#) map command. When multiple matches occur for the traffic, the order of precedence will be used to arrive at the final action. The order of precedence is as follows:

1. 802.1p bits
2. DSCP

The service ingress policy with *policy-id* 1 is a system-defined policy applied to services when no other policy is explicitly specified. This policy cannot be modified or deleted.

The **no** version of this command restores the factory default settings when used on *policy-id* 1. The default service ingress policy defines one queue associated with the best effort (**be**) forwarding class, with CIR of 0 and PIR of line rate.

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

The **no sap-ingress *policy-id*** command deletes the service 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.

Parameters

policy-id

uniquely identifies the policy. The *policy-name* cannot be used to create the policy, but it can be used to reference a policy that already exists.

Values 1 to 65535, or *policy-name* (up to 64 characters)

create

keyword used to create a sap-ingress policy

scope**Syntax**

scope {**exclusive** | **template**}

no scope

Context

config>qos>sap-ingress

Description

This command configures the Service Ingress QoS policy scope as **exclusive** or **template**.

The policy's scope cannot be changed if the policy is applied to a service.

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

Default

template

Parameters**exclusive**

when the scope of a policy is defined as **exclusive**, the policy can only be applied to one SAP. If a policy with an exclusive scope is assigned to a second SAP, an error message is generated. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.

template

when the scope of a policy is defined as **template**, the policy can be applied to multiple SAPs on the router

default-fc**Syntax**

default-fc *fc-name*

no default-fc

Context

config>qos>sap-ingress

Description

This command configures the default forwarding class for the policy. In the event that an ingress packet does not match a higher priority (more explicit) classification command, the default forwarding class will be associated with the packet. Unless overridden by an explicit forwarding class classification rule, all packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class. Optionally, the default ingress enqueueing priority for the traffic can be overridden as well.

The default forwarding class for **default-fc** is best effort (be). The **default-fc** settings are displayed in the **show configuration** and save output regardless of inclusion of the **detail** keyword.

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

Default

be

Parameters

fc-name

specifies the forwarding class name for the queue. The value given for *fc-name* must be one of the predefined forwarding classes in the system.

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

default-priority

Syntax

default-priority {high | low}

no default-priority

Context

config>qos>sap-ingress

Description

This command configures the default enqueueing priority for all packets received on an ingress SAP using this policy.

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

Default

low

Parameters

high

setting the enqueueing parameter to **high** for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing; once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

low

setting the enqueueing parameter to **low** for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing; once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

fc**Syntax**

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

Context

config>qos>sap-ingress

Description

This command is used to create a class of the forwarding class *fc-name*.

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

Parameters*fc-name*

specifies the forwarding class name for the queue. The value given for *fc-name* must be one of the predefined forwarding classes in the system.

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

Default n/a

create

keyword used to create a forwarding class

dot1p**Syntax**

dot1p *dot1p-priority* [**fc** *fc-name*] [**priority** {**high** | **low**}]

no dot1p *dot1p-priority*

Context

config>qos>sap-ingress

Description

This command explicitly sets the forwarding class and/or enqueueing priority when a packet is marked with a *dot1p-priority* specified. Adding a dot1p rule on the policy forces packets that match the *dot1p-priority* specified to override the forwarding class and enqueueing priority based on the parameters included in

the dot1p rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing priority derived from earlier matches in the classification hierarchy.

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 eight Class-of-Service (CoS) values commonly used to map packets to per-hop Quality-of-Service (QoS) behavior.

For QinQ applications, the dot1p bits used for classification are from either the top or bottom tag based on the **match-qinq-dot1p** configuration. For more information, see the command description and the “QinQ Dot1p Match Behavior” section in the 7705 SAR Services Guide.

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

Parameters

dot1p-priority

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

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

Values 0 to 7

fc-name

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

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

Default n/a

priority

the **priority** keyword is used to override the default enqueueing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule, the enqueueing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and the priority is not explicitly defined in the rule, the enqueueing priority is inherited based on previous rule matches.

high

the **high** keyword is used in conjunction with the **priority** keyword. Setting the enqueueing parameter to **high** for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing; once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

low

the **low** keyword is used in conjunction with the **priority** keyword. Setting the enqueueing parameter to low for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queueing; once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

dscp**Syntax**

dscp *dscp-name* [*dscp-name*...(up to 8 max)] [**fc** *fc-name*] [**priority** {**high** | **low**}]

no dscp *dscp-name* [*dscp-name*...(up to 8 max)]

Context

config>qos>sap-ingress

Description

This command explicitly sets the forwarding class and/or enqueueing priority when a packet is marked with the DiffServ Code Point (DSCP) value contained in *dscp-name*. Adding a DSCP rule on the policy forces packets that match the specified DSCP value to override the forwarding class and enqueueing priority based on the parameters included in the DSCP rule.

When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing priority derived from earlier matches in the classification hierarchy.

The DSCP value (referred to by *dscp-name*) is derived from the most significant six bits in the IP header ToS byte field (DSCP bits). The six DSCP bits define 64 DSCP values used to map packets to per-hop QoS behavior.

A list of up to 8 DSCP names can be specified to assign DSCP values to an FC mapping for DSCP classification.

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

Parameters*dscp-name*

the *dscp-name* is a required parameter that specifies the unique IP header ToS byte DSCP bits value that will match the DSCP rule.

A maximum of 64 DSCP rules are allowed on a single policy. The specified name must exist as a *dscp-name*. [Table 45: Valid DSCP names](#) lists all the valid DSCP names.

fc-name

the value given for *fc-name* must be one of the predefined forwarding classes in the system. Specifying the *fc-name* is optional. When a packet matches the rule, the forwarding class is only overridden when the *fc-name* parameter is defined on the rule.

If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

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

Default inherit (when *fc-name* is not defined, the rule preserves the previous forwarding class of the packet)

priority

this keyword overrides the default enqueueing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule, the enqueueing priority is only overridden when the **priority** keyword is defined on the rule. If the packet matches and **priority** is not explicitly defined in the rule, the enqueueing priority is inherited based on previous rule matches.

Default inherit

high

this keyword is used in conjunction with the **priority** keyword. Setting the enqueueing parameter to **high** for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing; once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

Default high

low

this keyword is used in conjunction with the **priority** keyword. Setting the enqueueing parameter to **low** for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing; once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

Default low

policy-name

Syntax

policy-name *policy-name*

no policy-name

Context

config>qos>sap-ingress

Description

This command configures a policy name for the SAP ingress policy.

Parameters

policy-name

specifies the name for the policy, up to 64 characters

queue

Syntax

queue *queue-id* [*queue-type*] [*queue-mode*] [**create**]

no queue *queue-id*

Context

config>qos>sap-ingress

Description

This command enables the context to configure a service ingress policy queue.

Explicit definition of an ingress queue's hardware scheduler status is supported. A single ingress queue allows support for multiple forwarding classes.

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

The expedited hardware schedulers are used to enforce expedited access to internal switch fabric destinations. The hardware status of the queue must be defined at the time of queue creation within the policy.

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

Parameters

queue-id

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

Values 1 to 8

queue-type

the **expedite**, **best-effort**, and **auto-expedite** queue types are mutually exclusive. Each defines the method that the system uses to service the queue from a hardware perspective. A keyword must be specified at the time the queue is created in the service ingress policy. If an attempt is made to change the keyword after the queue is initially defined, an error is generated.

expedite

the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue

best-effort

the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue

auto-expedite

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

Values expedite, best-effort, auto-expedite

Default auto-expedite

queue-mode

specifies the mode in which the queue is operating, either **priority-mode** or **profile-mode**. A keyword must be specified at the time the queue is created in the service ingress policy. If an attempt is made to change the keyword after the queue is initially defined, an error is generated.

Values priority-mode or profile-mode

Default priority-mode

priority-mode

configures the queue to be capable of handling traffic with two distinct priorities, high and low. These priorities are assigned by the stages preceding the queuing framework in the system. In priority mode, a packet's in-profile or out-of-profile state is determined by the state of the queue at scheduling time. When the queue rate is lower than or equal to the configured CIR, the packet is considered in-profile. When the queue rate is higher than the CIR, the packet is considered out-of-profile.

profile-mode

configures the queue to support color-aware profiling of the forwarding class mapped to the queue. Color-aware operational behavior is as follows.

- Forwarding classes defined as in-profile are handled as high priority and packets assigned to these forwarding classes are marked as in-profile. These profiled in-profile packets will consume queue CIR bandwidth.
- Forwarding classes defined as out-of-profile are handled as low priority and packets assigned to these forwarding classes are marked as out-of-profile. These profiled out-of-profile packets do not consume queue CIR for profile marking calculations.
- Forwarding classes that are not profiled (profile not set to in-profile or out-of-profile) are handled as high priority and are marked as in-profile or out-of-profile based on the dynamic rate of the ingress queue relative to its CIR. Non-profiled packets scheduled at or lower than the CIR will consume queue CIR bandwidth.

create

keyword used to create a sap-ingress queue context

6.4.2.1.8 Service ingress QoS policy forwarding class commands

broadcast-queue

Syntax

broadcast-queue *queue-id*

no broadcast-queue

Context

config>qos>sap-ingress>fc

Description

This command maps the broadcast forwarding type queue to the **fc** *fc-name*. The specified *queue-id* must already have been created within the policy before the mapping can be made. Once the forwarding class mapping is executed, all broadcast traffic on a SAP using this policy will be forwarded using the *queue-id*.

The **no** form of the command sets the broadcast forwarding type *queue-id* back to the default of no mapping to an FC.

Default

no broadcast-queue

Parameters

queue-id

an existing queue defined in the **config>qos>sap-ingress** context.

Values 1 to 8

de-1-out-profile

Syntax

[no] de-1-out-profile

Context

config>qos>sap-ingress>fc

Description

This command, when enabled on a parent forwarding class, applies a color profile mode to the packets stored in the queue associated with this forwarding class.

When this QoS policy is applied to the ingress of an FR VLL SAP, DE=1 frames are classified as out-of-profile and are not subject to the CIR marking.

All received DE=0 frames that exceed the CIR are marked as out-of-profile and have the DE set to 1, regardless of whether this command is enabled or disabled.

The **priority** option, if used, has no effect. All FR VLL DE=1 frames have their priority automatically set to low; DE=0 frames have their priority set to high.

All other capabilities of the Fpipe service are maintained. These capabilities include re-marking of the DE bit on egress SAP, and FR PW control word on egress network port for the packets that were classified as out-of-profile at the ingress SAP.

The **de-1-out-profile** command has an effect only when it is applied to the ingress of a SAP that is part of an Fpipe service.

The **no** form of the command disables the color profile mode of operation on all SAPs to which this ingress QoS policy is applied.

Default

no de-1-out-profile

multicast-queue

Syntax

multicast-queue *queue-id*

no multicast-queue

Context

config>qos>sap-ingress>fc

Description

This command maps the multicast forwarding type queue to the **fc** *fc-name*. The specified *queue-id* must already have been created within the policy before the mapping can be made. Once the forwarding class mapping is executed, all multicast traffic on a SAP using this policy will be forwarded using the *queue-id*.

The **no** form of the command sets the multicast forwarding type *queue-id* back to the default of no mapping to an FC.

Default

no multicast-queue

Parameters

queue-id

an existing queue defined in the **config>qos>sap-ingress** context.

Values 1 to 8

profile

Syntax

profile {in | out}

no profile

Context

config>qos>sap-ingress>fc

Description

This command specifies that ingress access data packets mapped to this forwarding class are either in-profile or out-of-profile.

Usually, at SAP ingress an access packet's state is defined as in-profile or out-of-profile based on the dynamic rate of the ingress queue before being forwarded to the fabric. Explicitly defining a forwarding class as in-profile or out-of-profile overrides this function by handling each packet according to the defined profile state; however, the forwarding class must be mapped to a queue configured for **profile-mode**. If the forwarding class is mapped to a queue configured for **priority-mode**, the forwarding class profile setting is ignored and the packet's state is defined as in-profile or out-of-profile based on the dynamic rate of the ingress queue.

When the **profile in** command is executed on a forwarding class that is mapped to a queue operating in profile mode, all packets associated with the class are handled as in-profile. When the **profile out** command is executed on a forwarding class that is mapped to a queue operating in profile mode, all packets associated with the class are handled as out-of-profile.

When the **no profile** command is executed on a forwarding class that is mapped to a queue operating in profile mode, the data packets using the forwarding class are marked as in-profile or out-of-profile based on the dynamic rate of the ingress queue relative to its CIR.

Packets explicitly handled as in-profile or out-of-profile flow through the same ingress service queue associated with the forwarding class to preserve order within flows. In-profile packets count against the CIR of the queue, reducing the amount of CIR available to forwarding classes using the queue that are not configured with an explicit profile. Out-of-profile packets do not count against the CIR of the queue, allowing classes using the queue that are not configured with an explicit profile to be measured against the full CIR.

A queue operating in profile mode can support in-profile, out-of-profile, and non-profiled packets simultaneously because multiple forwarding classes with different forwarding class profiles can be assigned to a single queue. In addition, a queue operating in profile mode is classified as high-priority or low-priority based on the configuration of the forwarding class profile, not the high-priority or low-priority configuration specified for DSCP or dot1p. Packets mapped to a forwarding class configured with **profile in** or **no profile** are marked as high priority. Packets mapped to a forwarding class configured with **profile out** are marked as low priority.

The **no** form of this command removes the explicit in-profile or out-of-profile configuration on a forwarding class.

Default

no profile

Parameters

- in

specifies that all packets associated with the forwarding class will be handled as in-profile
- out

specifies that all packets associated with the forwarding class will be handled as out-of-profile

queue

Syntax

- queue *queue-id*
- no queue

Context

config>qos>sap-ingress>fc

Description

This command overrides the default forwarding type queue mapping for **fc** *fc-name*.
The **no** form of this command sets the *queue-id* back to the default queue for the forwarding class (queue 1).

Parameters

- queue-id*

an existing queue defined in the **config>qos>sap-ingress** context
- Values

1 to 8
- Default

1

unknown-queue

Syntax

- unknown-queue *queue-id*
- no unknown-queue

Context

config>qos>sap-ingress>fc

Description

This command maps the unknown forwarding type queue to the **fc** *fc-name*. The specified *queue-id* must already have been created within the policy before the mapping can be made. Once the forwarding class

mapping is executed, all unknown forwarding type traffic on a SAP using this policy will be forwarded using the *queue-id*.

The **no** form of the command sets the unknown forwarding type *queue-id* back to the default of no mapping to an FC.

Default

no unknown-queue

Parameters

queue-id

an existing queue defined in the **config>qos>sap-ingress** context.

Values 1 to 8

6.4.2.1.9 Service queue QoS policy commands

adaptation-rule

Syntax

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

no adaptation-rule

Context

config>qos>sap-ingress>queue

config>qos>sap-egress>queue

Description

This command can be used to define how an operational rate is selected based on the configured PIR or CIR rate. Operational rates are the finite set of rates at which the schedulers on the network processor can operate.

The **no** form of the command removes any **adaptation-rule** constraints used to derive the operational rates for the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default

pir closest cir closest

Parameters

pir

pir defines the constraints enforced when adapting the PIR rate defined within the **queue queue-id rate** command. The **pir** keyword requires a qualifier that defines the constraint used when deriving the operational PIR rate for the queue. When the **rate** command is not specified, the default applies.

cir

cir defines the constraints enforced when adapting the CIR rate defined within the **queue queue-id rate** command. The **cir** keyword requires a qualifier that defines the constraint used when deriving the operational CIR rate for the queue. When the **cir** keyword is not specified, the default constraint applies.

adaptation-rule

specifies the constraints to be used while computing the operational CIR or PIR rate. The **max** (maximum), **min** (minimum), and **closest** parameters are mutually exclusive.

Values

max – causes the network processor to be programmed at an operational rate that is less than the configured PIR or CIR rate by up to 1.0%. For a network processor on a Gen-3 adapter card or platform, the average difference between the operational and the configured CIR rate is 2.0% (for frame sizes less than 2049 bytes) or 4.0% (for other frame sizes).

min – causes the network processor to be programmed at an operational rate that is greater than the configured PIR or CIR rate by up to 1.0%. For a network processor on a Gen-3 adapter card or platform, the average difference between the operational and the configured CIR rate is 2.0% (for frame sizes less than 2049 bytes) or 4.0% (for other frame sizes).

closest – causes the network processor to be programmed at an operational rate that is closest to the configured PIR or CIR rate

cbs**Syntax**

cbs {*size-in-kbytes* | **default**}

no cbs

Context

config>qos>sap-ingress>queue

config>qos>sap-egress>queue

Description

This command overrides the default committed buffer space (CBS) reserved for buffers of a specified queue. The value is configured in kilobytes.

The value in kilobytes is converted automatically to the number of buffers. The conversion calculation uses a non-configurable buffer size of 2304 bytes or 512 bytes, depending on the type of adapter card. See [Table 4: Buffer support on adapter cards and platforms](#) for a list of adapter cards and their associated buffers. The calculation is:

Number of buffers = Configured CBS value in bytes / Buffer size in bytes

At the egress of an $N > 1$ Apipe, the CBS value in a SAP egress QoS policy that is assigned to a SAP aggregation group causes n times that value of buffers to be committed, where n is the number of SAPs in

the SAP aggregation group. See the **show pools** command for information about how to view buffer pools for SAPs that are members of a SAP aggregation group. See the 7705 SAR Services Guide for information about how to configure SAP aggregation groups.

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

Default

"default" (8 kB for adapter cards and platforms with 512 byte buffer size) (18 kB for adapter cards and platforms with 2304 byte buffer size).

Parameters

size-in-kbytes
this parameter is an integer expression of the number of kilobytes reserved for the queue. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 to 131072

default
returns the CBS size to the default value

high-prio-only

Syntax

high-prio-only *percent*
no high-prio-only

Context

config>qos>sap-ingress>queue
config>qos>sap-egress>queue

Description

This command configures the percentage of buffer space for the queue, used exclusively by high-priority packets. The specified value overrides the default value for the context.

The priority of a packet can only be set in the service ingress policy and is only applicable on the ingress queues for a SAP. The profile state is used for enqueueing priority at sap-egress.

The **no** form of this command restores the default high-priority reserved size.

Parameters

percent
the percentage reserved for high priority traffic on the queue

Values 0 to 100 | default1

Default 10

mbs

Syntax

mbs *size* [bytes | kilobytes]

no mbs

Context

config>qos>sap-ingress>queue

config>qos>sap-egress>queue

Description

This command sets the maximum burst size (MBS) value for the buffers of the specified queue. The value is configured in bytes or kilobytes, and overrides the default MBS value. The default configuration is in kilobytes.

The **config>qos>sap-ingress>info detail** and **sap-egress>info detail** screens show the MBS in terms of bytes, unless it is a multiple of 1000. In that case, the display shows the MBS in kilobytes. For example, entering **mbs 200** or **mbs 200 kilobytes** configures and displays "200 kilobytes", entering **mbs 200000 bytes** also configures and displays "200 kilobytes", and entering **mbs 200100 bytes** configures and displays "200100 bytes".



Note: For the 7705 SAR, 1 kB of buffer management space is 1000 bytes.

The MBS value in bytes is converted automatically to the number of buffers. The conversion calculation uses a non-configurable buffer size of 2304 bytes or 512 bytes, depending on the type of adapter card. See [Table 4: Buffer support on adapter cards and platforms](#) for a list of adapter cards and their associated buffers. The calculation is:

Number of buffers = Configured MBS value in bytes / Buffer size in bytes (2304 or 512)

At the egress of an N > 1 Apipe, the MBS value in a SAP egress QoS policy that is assigned to a SAP aggregation group is used for each of the per-SAP queues for SAPs that are members of a SAP aggregation group. See the **show pools** command for information about how to view buffer pools for SAPs that are members of a SAP aggregation group. See the 7705 SAR Services Guide for information about how to configure SAP aggregation groups.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an adapter card can exceed the total amount of buffering available. Therefore, for a packet arriving at a queue that has not exceeded its MBS size, it is not guaranteed that a buffer will be available. If a buffer is not available, the packet will be discarded. RED/WRED slope parameters can be configured to control congestion in the case where the buffer capacity of the card is becoming exhausted.

Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard against queue starvation (that is, when a queue does not receive its fair share of buffers). Another safeguard is to properly set the RED/WRED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS size assigned to the queue to the default value.

Default

180 (kB) (converted to 78 packets when buffer size is 2304 bytes and to 351 packets when buffer size is 512 bytes)

Parameters

size

the *size* parameter is an integer expression of the maximum number of bytes of buffering allowed for the queue. A value of 4000 bytes or less causes the queue to discard all packets. Selecting **default** returns the MBS to the default value.

Values 0 to 131072000 | **default**

bytes

specifies that the size entered is in bytes

kilobytes

specifies that the size entered is in kB

rate

Syntax

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

no rate

Context

config>qos>sap-ingress>queue

config>qos>sap-egress>queue

Description

This command defines the administrative PIR and the administrative CIR parameters for the queue. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. Similarly, defining a CIR does not necessarily guarantee that the queue can schedule at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available bandwidth.

The **cir** keyword defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.

For service ingress, the PIR defines the maximum rate that the queue can transmit packets toward the fabric. The **cir** keyword defines the rate that packets are considered in-profile by the system. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

For service egress queues, the PIR defines the maximum rate that the queue can transmit packets out an egress interface.

When the PIR is set to **max** on a SAP-ingress queue, the max value defaults to the physical port line rate. On a SAP-egress queue, the PIR is set to the physical port line rate.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (max, 0).



Caution: The **rate** command can be executed at any time but should be executed during a maintenance window because the command can be service-affecting. Altering the PIR and CIR rates affects all queues created through the association of the service ingress or service egress QoS policy with the *queue-id*.



Note: The ingress traffic to an Epipe, Ipipe, IES, VPLS, and VPRN service may be shaped to a lower rate than the PIR and CIR values configured in the SAP ingress policy:

- at SAP ingress, the CIR and PIR settings under **rate** include both payload (customer) traffic and overhead traffic, which affects the shaping rate internally. Additional overhead bytes include the internal fabric header minus any bytes that have been removed from the original packet (such as the four-byte FCS).
- at SAP ingress, the actual shaping rate is related to the service rate (PIR or CIR) specified in the SAP ingress QoS policy, as shown below:

$$\text{shaping rate (actual)} = (\text{PIR or CIR}) / \text{ratio}$$

$$\text{where ratio} = (\text{customer packet size} + \text{additional bytes} - \text{removed headers} - 4 \text{ byte FCS}) / (\text{customer packet size})$$

- at SAP egress, shaping does not include the FCS, so the actual shaping rate is a bit higher than the PIR/CIR ratio configured in the QoS policy, as shown below:

$$\text{shaping rate (actual)} = (\text{PIR or CIR}) / \text{ratio}$$

$$\text{where ratio} = (\text{customer packet size} - 4 \text{ byte FCS}) / (\text{customer packet size})$$

Default

rate max **cir** 0 (this default specifies the amount of bandwidth in kb/s. The max value and the *pir-rate* value are mutually exclusive.)

Parameters

pir-rate

defines the administrative PIR rate, in kb/s, for the queue. When the rate command is executed, a valid PIR setting must be explicitly defined. When the rate command has not been executed, the default PIR of max is assumed. Fractional values are not allowed and must be given as a positive integer. The PIR rate has a minimum value of 8 kb/s.

The actual PIR rate is dependent on the queue's [adaptation-rule](#) parameters and the actual hardware where the queue is provisioned.

Values 1 to 100000000 | **max**

Default max



Note: If a PIR rate lower than 8 kb/s is specified, it is rounded up to this minimum value.

cir-rate

overrides the default administrative CIR used by the queue. When the rate command is executed, a *cir-rate* setting is optional. When the rate command has not been executed or

the **cir** keyword is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer. The CIR rate has a minimum value of 8kb/s.

Values 0 to 100000000 | **max**

Default 0



Note: If a CIR rate lower than 8 kb/s is specified, it is rounded up to this minimum value (with the exception of 0, which does not get rounded up).

slope-policy

Syntax

slope-policy *name*
no slope-policy

Context

config>qos>sap-ingress queue
config>qos>sap-egress queue

Description

This command specifies the slope parameters controlling the queue.

Default

slope-policy default

Parameters

name
the name of the slope policy

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

6.4.2.2 Operational commands

copy

Syntax

copy sap-egress *src-pol dst-pol* [**overwrite**]

copy sap-ingress *src-pol dst-pol* [**overwrite**]

Context

config>qos

Description

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

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

Parameters

src-pol dst-pol

indicates that the source policy ID and the destination policy ID are sap-egress or 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 that the existing destination policy is to be replaced. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, an error will occur if the destination policy ID exists.

6.4.2.3 Show commands



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

sap-egress

Syntax

sap-egress [*policy-id*] [**standard** | **mc-mlppp**] [**association** | **match-criteria** | **detail**]

sap-egress summary

Context

show>qos

Description

This command displays service egress and MC-MLPPP SAP egress QoS policy information.

Parameters

policy-id

displays information about the specific policy ID

Values 1 to 65535, or *policy-name* (up to 64 characters)

Default all service egress policies

- standard**
displays all standard SAP egress QoS policies
- mc-mlppp**
displays only MC-MLPPP SAP egress policy information
- association**
displays policy associations
- match-criteria**
displays match criteria
- detail**
displays detailed policy information including policy associations
- summary**
displays a summary of the number of SAP egress policies

Output

The following output is an example of service egress and MC-MLPPP SAP egress QoS policy information, and [Table 46: SAP egress field descriptions](#) describes the fields.

Output example

```
*A:ALU-1# show qos sap-egress
=====
Sap Egress Policies
=====
Policy-Id   Type      Scope    Name      Description
-----
1           Default   Template default   Default SAP egress QoS policy.
200         Standard Template
300         MC-MLPPP Template
400         Standard Exclusive
500         Standard Template
550         MC-MLPPP Template
600         Standard Exclusive
-----
Number of Policies : 7
-----
*A:ALU-1#

*A:Sar18 Dut-B>show>qos# sap-egress summary
=====
Sap Egress Summary
=====
Number of Policies      : 7
Active Policies        : 6
Inactive Policies       : 1
=====
```

```
*A: Sar18 Dut-B>show>qos#
```

```
*A: ALU-1# show qos sap-egress 500 detail
```

```
=====
QoS Sap Egress
=====
```

```
-----
Standard Sap Egress Policy (500)
-----
```

```
Policy-id      : 500                      Scope      : Template
Name           : sap_egress_policy_500
Description    : Egress Policy 500
Pkt.Byte Offset: add 0
Policy Active  : False
```

```
-----
Queue Information
-----
```

```
Queue-Id      : 1                      Queue-Type   : auto-expedite
Admin PIR     : max                    Admin CIR    : 0
PIR Rule      : closest                CIR Rule     : closest
CBS           : Def                   MBS          : Def
Slope Policy  : default                Packet Byte Offset : def
High Priority  : def

Queue-Id      : 4                      Queue-Type   : auto-expedite
Admin PIR     : max                    Admin CIR    : 0
PIR Rule      : closest                CIR Rule     : closest
CBS           : 2 KB                  MBS          : Def
Slope Policy  : default                Packet Byte Offset : def
High Priority  : def
```

```
-----
FC   Queue Dot1p Exp/Def      DSCP Marking
-----
be   def   Default            default
ef   def   Default            DSCP In:cp2 Out:cp3
```

```
-----
Associations
-----
```

```
No Associations Found.
```

```
=====
*A: ALU-1#
```

```
*A: 7705: Dut-C# show qos sap-egress mc-mlppp 550 detail
```

```
=====
QoS Sap Egress
=====
```

```
-----
MC-MLPPP Sap Egress Policy (550)
-----
```

```
Policy-id      : 550                      Scope      : Template
Name           : (Not Specified)
Description    : New Multi-class type policy
Pkt.Byte Offset: default
Policy Active  : False
```

```
-----
Queue Information
-----
```

```
Queue-Id      : 1                      Queue-Type   : auto-expedite
Admin PIR     : max
PIR Rule      : closest
```

```

CBS           : Def           MBS           : Def
Slope Policy  : default
High Priority  : def

```

```

-----
FC Name  Queue-id  DSCP Marking
-----
No Matching Entries

-----
Associations
-----
No Associations Found.

```

Table 46: SAP egress field descriptions

Label	Description
Summary	
Number of Policies	The number of SAP egress policies
Active Policies	The number of active SAP egress policies
Inactive Policies	The number of inactive SAP egress policies
QoS Sap Egress Policy	
Policy-Id	The ID that uniquely identifies the policy
Type	Indicates the type of SAP egress policy; Default, Standard, or MC-MLPPP
Scope	Exclusive - this policy can only be applied to a single SAP
	Template - this policy can be applied to multiple SAPs on the router
Name	The policy name
Description	A text string that helps identify the policy's context in the configuration file
Pkt.Byte Offset	Indicates the value of the packet byte offset applied to the packet for scheduling, if applicable. A value of "default" indicates that legacy mode packet scheduling is in use, in which packets are scheduled based on size including internal overhead.
Policy Active	Indicates whether the policy is active or inactive
Queue Information	
Queue-Id	The ID that identifies the queue

Label	Description
Queue-Type	The queue type: auto-expedite, expedite, or best-effort
Admin PIR	The administrative Peak Information Rate (PIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress interface.
Admin CIR	The administrative Committed Information Rate (CIR) parameters for the queue. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.
PIR Rule	min - the operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command
	max - the operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command
	closest - the operational PIR for the queue will be the rate closest to the rate specified using the rate command
CIR Rule	min - the operational CIR for the queue will be equal to or greater than the administrative rate specified using the rate command except where the derived operational CIR is greater than the operational PIR. If the derived operational CIR is greater than the derived operational PIR, the operational CIR will be made equal to the operational PIR.
	max - the operational CIR for the queue will be equal to or less than the administrative rate specified using the rate command
	closest - the operational CIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational PIR
CBS	def - the CBS value reserved for the queue
	value - the value to override the default reserved buffers for the queue
MBS	def - the MBS value reserved for the queue
	value - the value to override the default maximum size for the queue
Slope Policy	The slope policy for the queue
Packet Byte Offset	<p>The value of the packet byte offset applied to the queue.</p> <p>A value of "default" indicates that legacy mode packet scheduling is in use, in which packets are scheduled based on size including internal overhead.</p>

Label	Description
High Priority	The percentage of buffer space for the queue, used exclusively by high-priority packets
FC Mapping Information	
FC	The forwarding class
Queue	The queue ID that uniquely identifies the queue within the policy
Dot1p Exp/Def (Explicit/Default)	The egress dot1p bits used for the FC-to-dot1p mapping: Explicit (an egress dot1p bit is configured for the FC-to-egress-dot1p bit mapping) or Default
DSCP Marking	The DSCP name and DSCP priority bits mapping for the forwarding class
Associations	
Associations	Service Name - the name of the service
	Service-Id - the unique service ID number that identifies the service in the service domain
	Customer-Id - the customer ID that identifies the customer to the service
	SAP - the Service Access Point within the service where the policy is applied

sap-ingress

Syntax

sap-ingress [*policy-id*] [**association** | **match-criteria** | **detail**]

sap-ingress summary

Context

show>qos

Description

This command displays service ingress QoS policy information.

Parameters

policy-id

displays information about the specific policy

Values 1 to 65535, or *policy-name* (up to 64 characters)

Default all service ingress policies

association

displays policy associations

match-criteria

displays match criteria

detail

displays detailed policy information including policy associations

summary

displays a summary of the number of SAP ingress policies

Output

The following output is an example of service ingress QoS policy information, and [Table 47: SAP ingress field descriptions](#) describes the fields.

Output example

```
A:ALU-1# show qos sap-ingress
=====
Sap Ingress Policies
=====
Policy-Id  Scope    Name          Description
-----
1          Template default      Default SAP ingress QoS policy.
100        Template
=====
Number of Policies : 2
=====
*A:ALU-1#
```

```
*A:Sar18 Dut-B>show>qos# sap-ingress summary
=====
Sap Ingress Summary
=====
Number of Policies      : 2
Active Policies         : 1
Inactive Policies       : 1
=====
*A:Sar18 Dut-B>show>qos#
```

```
A:ALU-1# show qos sap-ingress 1 detail
=====
QoS Sap Ingress
=====
Sap Ingress Policy (1)
-----
Policy-id      : 1                Scope      : Template
Default FC     : be              Priority     : Low
Name           : default
Description     : Default SAP ingress QoS policy.
Pkt.Byte Offset: default
Policy Active  : True
-----
```



```

Q  Mode      CIR Admin PIR Admin CBS      HiPrio Packet Slope-Policy
      CIR Rule PIR Rule MBS      Offset
-----
1  Prio      0          max      def      def      default default
      closest closest  def
-----
FC          UCastQ      MCastQ      BCastQ      UnknownQ
-----
No FC-Map Entries Found.
-----
FC          DE-1-out-profile Profile
-----
No FC-Map Entries Found.
-----
Dot1p      FC          Priority
-----
No Dot1p-Map Entries Found.
-----
DSCP      FC          Priority
-----
No DSCP-Map Entries Found.
-----
Associations
-----
Service-Name : epipel
Service-Id   : 1 (Epipe)      Customer-Id : 1
- SAP : 1/1/2
- SAP : 1/12/3:100.100
Service-Id   : 4 (Cpipe)      Customer-Id : 1
- SAP : 1/2/1.1
Service-Id   : 5 (IES)        Customer-Id : 1
- SAP : 1/10/4
Service-Name : VPRN21_hybrid
Service-Id   : 21 (VPRN)      Customer-Id : 1
- SAP : 1/7/6
=====
*A:7705custDoc:Sar18#

```

```
*A:7705:Dut-A# show qos sap-ingress 2 detail
```

```
=====
QoS Sap Ingress
=====
```

```
Sap Ingress Policy (2)
```

```

Policy-id      : 2          Scope      : Template
Default FC     : be        Priority    : Low
Name           : (Not Specified)
Description     : (Not Specified)
Pkt.Byte Offset: default
Policy Active   : True
-----

```

```

Q  Mode      CIR Admin PIR Admin CBS      HiPrio Packet Slope-Policy
      CIR Rule PIR Rule MBS      Offset
-----
1  Prio      0          max      def      def      def      default
      closest closest  def
2  Prio      0          max      def      def      def      default

```

```

3 Profile 0 closest closest def
          closest max def
          closest def def      def      def      default

-----
FC          UCastQ          MCastQ          BCastQ          UnknownQ
-----
nc          3          def          def          def
-----
FC          DE-1-out-profile Profile
-----
nc          No          In

-----
Dot1p          FC          Priority
-----
No Dot1p-Map Entries Found.

-----
DSCP          FC          Priority
-----
No DSCP-Map Entries Found.

-----
Associations
-----
Service-Name   : XYZ Epipe 1301
Service-Id     : 1301 (Epipe)          Customer-Id   : 1
- SAP : 1/2/8:1

=====

```

A:ALU-1# show qos sap-ingress 100 detail

=====

QoS Sap Ingress

=====

Sap Ingress Policy (100)

```

-----
Policy-id      : 100          Scope      : Template
Default FC    : be          Priority    : Low
Name          : (Not Specified)
Description    : Used on VPN SAP
Pkt.Byte Offset: default
Policy Active  : True
-----

```

```

-----
Q  Mode      CIR Admin PIR Admin CBS      HiPrio  Packet  Slope-Policy
      CIR Rule PIR Rule MBS      Offset
-----
1  Prio      0          max      def      def      default default
      closest closest def
2  Prio      25000     2500     1500     10      default default
      closest closest 10
-----

```

```

-----
FC          UCastQ          MCastQ          BCastQ          UnknownQ
-----
h2          def          def          def          def
ef          2          def          def          def
h1          2          def          def          def

```

```

-----
FC          DE-1-out-profile  Profile
-----
No FC-Map Entries Found.

-----
Dot1p      FC                      Priority
-----
1          be                      Low

-----
DSCP       FC                      Priority
-----
be         be                      Low

-----
Associations
-----
No Associations Found.
=====
*A:ALU-1#

```

Table 47: SAP ingress field descriptions

Label	Description
Summary	
Number of Policies	The number of SAP ingress policies
Active Policies	The number of active SAP ingress policies
Inactive Policies	The number of inactive SAP ingress policies
QoS Sap Ingress Policy	
Policy-Id	The ID that uniquely identifies the policy
Scope	Exclusive - this policy can only be applied to a single SAP
	Template - this policy can be applied to multiple SAPs on the router
Default FC	The default forwarding class for the policy
Priority	The default enqueueing priority
Name	The policy name
Description	A text string that helps identify the policy's context in the configuration file
Pkt.Byte Offset	Indicates the value of the packet byte offset applied to the packet for scheduling, if applicable. A value of "default" indicates that legacy mode packet scheduling is in use, in which packets are scheduled based on size including internal overhead.

Label	Description
Policy Active	Indicates whether the policy is active or inactive
Queue Information	
Q	The queue number
Mode	The mode for the queue, either priority mode or profile mode
CIR Admin	The 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 Rule	min - the operational CIR for the queue will be equal to or greater than the administrative rate specified using the rate command except where the derived operational CIR is greater than the operational PIR. If the derived operational CIR is greater than the derived operational PIR, the operational CIR will be made equal to the operational PIR.
	max - the operational CIR for the queue will be equal to or less than the administrative rate specified using the rate command
	closest - the operational CIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational PIR
PIR Admin	The administrative PIR parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric.
PIR Rule	min - the operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command
	max - the operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command
	closest - the operational PIR for the queue will be the rate closest to the rate specified using the rate command
CBS	def - the CBS value reserved for the queue
	value - the value to override the default reserved buffers for the queue
MBS	def - the MBS value reserved for the queue
	value - the value to override the default maximum size for the queue
HiPrio	The percentage of buffer space for the queue, used exclusively by high-priority packets

Label	Description
Packet Offset	The value of the packet byte offset applied to the queue. A value of "default" indicates that legacy mode packet scheduling is in use, in which packets are scheduled based on size including internal overhead.
Slope-Policy	The slope policy for the queue
Forwarding Class Information	
UCastQ	The unicast forwarding type queue mapping (default or queue number)
MCastQ	The multicast forwarding type queue mapping (default or queue number)
BCastQ	The broadcast forwarding type queue mapping (default or queue number)
UnknownQ	The unknown unicast forwarding type queue mapping (default or queue number)
Profile	The in-profile or out-of-profile state for packets assigned to this forwarding class
FC DE-1-out-profile Information	
DE-1-out-profile	When enabled on a Frame Relay VLL, DE-1 frames are classified as out-of-profile and are not subject to the CIR marking
Dot1p Information	
Dot1p	The forwarding class and/or enqueueing priority when a packet is marked with a <i>dot1p-value</i> specified
FC	The specified forwarding class used for the dot1p-to-FC mapping when a packet matches the dot1p rule. If no FC is specified, the default FC is used.
Priority	The optional priority setting overrides the default enqueueing priority for the packets received on an ingress SAP that uses the policy that matches this rule. The FC priority setting only applies to an FC that maps to a priority mode queue.
	High - the high enqueueing parameter for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. This parameter only applies to an FC that maps to a priority mode queue.
	Low - the low enqueueing parameter for a packet decreases the likelihood of enqueueing the packet when the ingress queue is

Label	Description
	congested. This parameter only applies to an FC that maps to a priority mode queue.
DSCP Information	
DSCP	The forwarding class and/or enqueueing priority when a packet is marked with the DiffServ Code Point (DSCP) value.
FC	The specified forwarding class used for the DSCP-to-FC mapping when a packet matches the DSCP rule. If no FC is specified, the default FC is used.
Priority	The default enqueueing priority overrides for all packets received on an ingress SAP using this policy that match this rule. The FC priority setting only applies to an FC that maps to a priority mode queue.
	High - the high enqueueing parameter for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. This parameter only applies to an FC that maps to a priority mode queue.
	Low - the low enqueueing parameter for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. This parameter only applies to an FC that maps to a priority mode queue.
Associations	
Associations	Service Name - the name of the service
	Service-Id - the unique service ID number that identifies the service in the service domain
	Customer-Id - the customer ID that identifies the customer to the service
	SAP - the SAP within the service where the service ingress policy is applied

pools

Syntax

pools mda mda-id [detail]

pools mda mda-id [egress | ingress] [ring | v-port] [detail]

Context

show

Description

This command displays buffer pool information for an adapter card. This information pertains to the memory allocation that is used for queuing purposes. The information is displayed according to the number of allocated buffers.

Because the 10-port 1GigE/1-port 10GigE X-Adapter card has separate ingress and egress buffer pools, there are two sets of pool information displayed for the card. Use the **egress** keyword to display only egress buffer pool information; use the **ingress** keyword to display only ingress buffer pool information; use the **mda** keyword by itself to display both egress buffer pool information and ingress buffer pool information.



Note: The **egress** and **ingress** keywords only apply to the 10-port 1GigE/1-port 10GigE X-Adapter card. The **ring** and **v-port** keywords only apply to the 2-port 10GigE (Ethernet) Adapter card and 2-port 10GigE (Ethernet) module.

All adapter cards allocate a fixed-size space for each buffer. The 7705 SAR supports two buffer sizes: 2304 and 512 bytes, depending on the type of adapter card. See [Table 4: Buffer support on adapter cards and platforms](#) for a list of adapter cards and their associated buffers.

A buffer size of 2304 bytes is large enough to accommodate the maximum frame size supported on the non-buffer-chaining adapter cards on the 7705 SAR. The adapter cards that support a buffer size of 2304 bytes always have a 1-to-1 correspondence of packets to buffers.

A buffer size of 512 bytes is smaller than the largest frame size supported on the 7705 SAR. Packets that are larger than 512 bytes use more than one buffer. The adapter cards that support a buffer size of 512 bytes dynamically divide packets that are larger than 512 bytes into a series of concatenated buffers in a method called buffer chaining. See [Buffer unit allocation and buffer chaining](#) for more information.

Adapter cards that support byte-based WRED use 512-byte buffers and can have a many-to-1 correspondence of packets to buffers. For example, assuming 128-byte packets, to fill 512 bytes in a queue takes four buffers using payload-based WRED (also called byte-based WRED), compared to one buffer in buffer-based WRED. When using byte-based WRED for frame sizes smaller than 512 bytes (128 or 256 bytes), the actual number of buffers used in the queue may exceed the MBS configured for the queue, since byte-based WRED does not count buffer overhead (filler) bytes in the queue to determine packet discard or tail drop eligibility, as is the case for buffer-based WRED. See [Payload-based WRED](#) for more information.

Parameters

mda-id

the location of the adapter card (in the form *slot/card*)

mda

displays buffer pool information for the adapter card

mda detail

displays detailed buffer pool information for the adapter card

egress

displays egress buffer pool information for the 10-port 1GigE/1-port 10GigE X-Adapter card

egress detail

displays detailed egress buffer pool information for the 10-port 1GigE/1-port 10GigE X-Adapter card

ingress

displays ingress buffer pool information for the 10-port 1GigE/1-port 10GigE X-Adapter card

ingress detail

displays detailed ingress buffer pool information for the 10-port 1GigE/1-port 10GigE X-Adapter card

ring

displays ring buffer pool information for the 2-port 10GigE (Ethernet) Adapter card and 2-port 10GigE (Ethernet) module

v-port

displays v-port buffer pool information for the 2-port 10GigE (Ethernet) Adapter card and 2-port 10GigE (Ethernet) module

Output

The following outputs are examples of detailed buffer pool information for an 8-port Gigabit Ethernet Adapter card, a 32-port T1/E1 ASAP Adapter card, a 10-port 1GigE/1-port 10GigE X-Adapter card, and a 2-port 10GigE (Ethernet) Adapter card. [Table 48: Buffer pool field descriptions](#) describes the fields.

- [Output example \(summary, 8-port Gigabit Ethernet Adapter card with a 512-byte buffer\)](#)
- [Output example \(summary, 32-port T1/E1 ASAP Adapter card with a 2304-byte buffer\)](#)

The following outputs are examples of ingress and egress buffer pool information for the 10-port 1GigE/1-port 10GigE X-Adapter card, and v-port and ring buffer pool information for the 2-port 10GigE (Ethernet) Adapter card. [Table 48: Buffer pool field descriptions](#) describes the fields.

- [Output example \(summary, 10-port 1GigE/1-port 10GigE X-Adapter card\)](#)
- [Output example \(summary, egress, 10-port 1GigE/1-port 10GigE X-Adapter card\)](#)
- [Output example \(detailed, egress, 10-port 1GigE/1-port 10GigE X-Adapter card\)](#)
- [Output example \(summary, ingress, 10-port 1GigE/1-port 10GigE X-Adapter card\)](#)
- [Output example \(detailed, ingress, 10-port 1GigE/1-port 10GigE X-Adapter card\)](#)
- [Output example \(2-port 10GigE \(Ethernet\) Adapter card\)](#)

Output example (summary, 8-port Gigabit Ethernet Adapter card with a 512-byte buffer)

The following output is an example of buffer pool information for an 8-port Gigabit Ethernet Adapter card. Outputs for the 6-port Ethernet 10Gbps Adapter card are similar. The buffer size (512 bytes) is used by all Ethernet adapter cards. [Table 48: Buffer pool field descriptions](#) describes the fields.

```
*A:ALU-1# show pools mda 1/1 detail
=====
Buffer Pool Information
=====
Pool Total       : 524287 buffers   Buffer Size      : 512 bytes
Pool Shared      : 387383 buffers   Pool Resv       : 136904 buffers
Pool Total In Use : 62 buffers
Pool Exhaustion Drop : 0
=====
=====
Access Ingress Queues
```


Name	FC-Maps	0.MBS(buf) 0.CBS(buf)	Depth(buf)
1->1/1/1:10->1	be l2 af l1 h2 ef h1 nc	351 16	1
1->1/1/1:20->1	be l2 af l1 h2 ef h1 nc	351 16	1

Access Egress Queues

Name	FC-Maps	0.MBS(buf) 0.CBS(buf)	Depth(buf)
1->1/1/1:10->1	be l2 af l1 h2 ef h1 nc	351 16	1
1->1/1/1:20->1	be l2 af l1 h2 ef h1 nc	351 16	1

Network Ingress Queues

FC-Maps	Dest	0.MBS(buf) 0.CBS(buf)	Depth(buf)
be	1/1	26214 528	0
l2	1/1	26214 1312	0
af	1/1	26214 3936	0
l1	1/1	13107 1312	0
h2	1/1	26214 3936	0
ef	1/1	26214 3936	0
h1	1/1	13107 1312	0
nc	1/1	13107 1312	0
be	1/2	26214 528	0
l2	1/2	26214 1312	0
af	1/2	26214 3936	0
l1	1/2	13107 1312	0
h2	1/2	26214 3936	0
ef	1/2	26214 3936	0
h1	1/2	13107	0

nc	1/2	1312	0
		13107	
be	1/3	1312	0
		26214	
l2	1/3	528	0
		26214	
af	1/3	1312	0
		26214	
l1	1/3	3936	0
		13107	
h2	1/3	1312	0
		26214	
ef	1/3	3936	0
		26214	
h1	1/3	3936	0
		13107	
nc	1/3	1312	0
		13107	
be	1/4	1312	0
		26214	
l2	1/4	528	0
		26214	
af	1/4	1312	0
		26214	
l1	1/4	3936	0
		13107	
h2	1/4	1312	0
		26214	
ef	1/4	3936	0
		26214	
h1	1/4	3936	0
		13107	
nc	1/4	1312	0
		13107	
be	1/5	1312	0
		26214	
l2	1/5	528	0
		26214	
af	1/5	1312	0
		26214	
l1	1/5	3936	0
		13107	
h2	1/5	1312	0
		26214	
ef	1/5	3936	0
		26214	
h1	1/5	3936	0
		13107	
nc	1/5	1312	0
		13107	
be	1/6	1312	0
		26214	
l2	1/6	528	0
		26214	
af	1/6	1312	0
		26214	
l1	1/6	3936	0
		13107	
h2	1/6	1312	0
		26214	
ef	1/6	3936	0
		26214	
h1	1/6	3936	0
		13107	

nc	1/6	1312	
		13107	0
be	MCast	1312	
		26214	0
l2	MCast	528	
		26214	0
af	MCast	528	
		26214	0
l1	MCast	528	
		13107	0
h2	MCast	528	
		26214	0
ef	MCast	528	
		26214	0
h1	MCast	528	
		13107	0
nc	MCast	528	
		13107	0
		528	
=====			
Network Egress Queues			
=====			
FC-Maps	ID	0.MBS(buf)	Depth(buf)
		0.CBS(buf)	

be	1/1/2	26214	0
		528	
l2	1/1/2	26214	0
		1312	
af	1/1/2	26214	0
		3936	
l1	1/1/2	13107	0
		1312	
h2	1/1/2	26214	0
		3936	
ef	1/1/2	26214	0
		3936	
h1	1/1/2	13107	0
		1312	
nc	1/1/2	13107	0
		1312	
be	1/1/1:100->interface1	26214	0
		528	
l2	1/1/1:100->interface1	26214	0
		1312	
af	1/1/1:100->interface1	26214	0
		3936	
l1	1/1/1:100->interface1	13107	0
		1312	
h2	1/1/1:8->test	26214	0
		3936	
ef	1/1/1:8->test	26214	0
		3936	
h1	1/1/1:8->test	13107	0
		1312	
nc	1/1/1:8->test	13107	0
		1312	
=====			
*A:ALU-1#			

Output example (summary, 32-port T1/E1 ASAP Adapter card with a 2304-byte buffer)

The following output is an example of buffer pool information for a 32-port T1/E1 ASAP Adapter card that supports an ATM N > 1 SAP aggregation group. The buffer size used by the adapter card is 2304 bytes, and the naming convention used at the access ingress queue identifies a SAP aggregation group for an N > 1 ATM pseudowire. [Table 48: Buffer pool field descriptions](#) describes the fields.

```
*A:ALU-1# show pools mda 1/1 detail
=====
Buffer Pool Information
=====
Pool Total      : 16520 buffers      Buffer Size      : 2304 bytes
Pool Shared     : 10776 buffers      Pool Resv       : 5744 buffers
Pool Total In Use : 489 buffers
Pool Exhaustion Drop : 0

=====

=====
Access Ingress Queues
=====
-----
Name                FC-Maps      0.MBS(buf) Depth(buf)
                  0.CBS(buf)
-----
10->SAG1->1
                  be l2 af l1   434      434
                  h2 ef h1 nc    8

=====

=====
Access Egress Queues
=====
-----
Name                FC-Maps      0.MBS(buf) Depth(buf)
                  0.CBS(buf)
-----
10->1/1/1.1:0/32->1
                  be l2 af l1   78        2
                  h2 ef h1 nc    8
10->1/1/2.1:0/33->1
                  be l2 af l1   78        1
                  h2 ef h1 nc    8
10->bundle-ima-1/1.1:0/34->1
                  be l2 af l1   78        1
                  h2 ef h1 nc    8

=====

=====
Network Ingress Queues
=====
-----
FC-Maps            Dest      0.MBS(buf) Depth(buf)
                  0.CBS(buf)
-----
No Match Found

=====

=====
Network Egress Queues
=====
-----
FC-Maps ID        0.MBS(buf) Depth(buf)
                  0.CBS(buf)
-----
```

```
No Match Found
```

```
*A:ALU-1#
```

Output example (summary, 10-port 1GigE/1-port 10GigE X-Adapter card)

The following output is an example of showing both ingress and egress buffer pool information for a 10-port 1GigE/1-port 10GigE X-Adapter card. [Table 48: Buffer pool field descriptions](#) describes the fields.

```
*A:ALU-1# show pools mda 1/X2
```

Ingress Buffer Pool Information

```
Pool Total      : 524287 buffers   Buffer Size      : 512 bytes
Pool Shared     : 228975 buffers   Pool Resv       : 295312 buffers
Pool Total In Use : 55 buffers
Pool Exhaustion Drop : 0
```

Egress Buffer Pool Information

```
Pool Total      : 524287 buffers   Buffer Size      : 512 bytes
Pool Shared     : 505415 buffers   Pool Resv       : 18872 buffers
Pool Total In Use : 6 buffers
Pool Exhaustion Drop : 0
```

```
*A:ALU-1#
```

Output example (summary, egress, 10-port 1GigE/1-port 10GigE X-Adapter card)

The following output is an example of showing only summary egress buffer pool information for a 10-port 1GigE/1-port 10GigE X-Adapter card. [Table 48: Buffer pool field descriptions](#) describes the fields.

```
*A:ALU-1# show pools mda 1/X2 egress
```

Egress Buffer Pool Information

```
Pool Total      : 524287 buffers   Buffer Size      : 512 bytes
Pool Shared     : 505415 buffers   Pool Resv       : 18872 buffers
Pool Total In Use : 6 buffers
Pool Exhaustion Drop : 0
```

Output example (detailed, egress, 10-port 1GigE/1-port 10GigE X-Adapter card)

The following output is an example of showing only detailed egress buffer pool information for a 10-port 1GigE/1-port 10GigE X-Adapter card. [Table 48: Buffer pool field descriptions](#) describes the fields.

```
*A:ALU-1# show pools mda 1/X2 egress detail
```

Egress Buffer Pool Information

```
Pool Total      : 524287 buffers   Buffer Size      : 512 bytes
Pool Shared     : 505415 buffers   Pool Resv       : 18872 buffers
Pool Total In Use : 6 buffers
Pool Exhaustion Drop : 0
```

```
=====
Access Egress Queues
=====
```

```
-----
Name                FC-Maps    0.MBS(buf) Depth(buf)
                  0.CBS(buf)
-----
```

```
No Match Found
=====
```

```
=====
Network Egress Queues
=====
```

```
-----
FC-Maps ID          0.MBS(buf) Depth(buf)
                  0.CBS(buf)
-----
```

```
be          1/X2/1          26214    0
                    528
l2          1/X2/1          26214    0
                    1312
af          1/X2/1          26214    0
                    3936
l1          1/X2/1          13107    0
                    1312
h2          1/X2/1          26214    0
                    3936
ef          1/X2/1          26214    0
                    3936
h1          1/X2/1          13107    0
                    1312
nc          1/X2/1          13107    0
                    1312
=====
```

```
=====
IPv4 GRE Fragment Reassembly Queue Groups
=====
```

```
-----
FC-Maps ID  Profile ID    0.MBS(buf) Depth(buf)
                  0.CBS(buf)
-----
```

```
be          1          26214    0
                    528
l2          1          26214    0
                    1312
af          1          26214    0
                    3936
l1          1          13107    0
                    1312
h2          1          26214    0
                    3936
ef          1          26214    0
                    3936
h1          1          13107    0
                    1312
nc          1          13107    0
                    1312
....
be          2          6214     0
                    528
```

```
l2                2                6214        0
                  312
. . . .
=====
```

Output example (summary, ingress, 10-port 1GigE/1-port 10GigE X-Adapter card)

The following output is an example of showing only summary ingress buffer pool information for a 10-port 1GigE/1-port 10GigE X-Adapter card. [Table 48: Buffer pool field descriptions](#) describes the fields.

```
*A:ALU-1# show pools mda 1/X2 ingress
=====
Ingress Buffer Pool Information
=====
Pool Total       : 524287 buffers   Buffer Size      : 512 bytes
Pool Shared      : 228975 buffers   Pool Resv        : 295312 buffers
Pool Total In Use : 55 buffers
Pool Exhaustion Drop : 0
=====
*A:ALU-1#
```

Output example (detailed, ingress, 10-port 1GigE/1-port 10GigE X-Adapter card)

The following output is an example of showing only detailed ingress buffer pool information for a 10-port 1GigE/1-port 10GigE X-Adapter card. [Table 48: Buffer pool field descriptions](#) describes the fields.

```
A:ALU-1# show pools mda 1/X2 ingress detail
=====
Ingress Buffer Pool Information
=====
Pool Total       : 524287 buffers   Buffer Size      : 512 bytes
Pool Shared      : 228975 buffers   Pool Resv        : 295312 buffers
Pool Total In Use : 55 buffers
Pool Exhaustion Drop : 0
=====

=====
Access Ingress Queues
=====
-----
Name                FC-Maps    0.MBS(buf) Depth(buf)
                  0.CBS(buf)
-----
No Match Found
=====

=====
Network Ingress Queues
=====
-----
FC-Maps             Dest      0.MBS(buf) Depth(buf)
                  0.CBS(buf)
-----
be                  1/1      26214      0
                  528
l2                  1/1      26214      0
                  1312
af                  1/1      26214      0
                  3936
l1                  1/1      13107      0
```

		1312	
h2	1/1	26214	0
		3936	
ef	1/1	26214	0
		3936	
h1	1/1	13107	0
		1312	
nc	1/1	13107	0
		1312	
be	1/2	26214	0
		528	
l2	1/2	26214	0
		1312	
af	1/2	26214	0
		3936	
l1	1/2	13107	0
		1312	
h2	1/2	26214	0
		3936	
ef	1/2	26214	0
		3936	
h1	1/2	13107	0
		1312	
nc	1/2	13107	0
		1312	
be	1/3	26214	0
		528	
l2	1/3	26214	0
		1312	
af	1/3	26214	0
		3936	
l1	1/3	13107	0
		1312	
h2	1/3	26214	0
		3936	
ef	1/3	26214	0
		3936	
h1	1/3	13107	0
		1312	
nc	1/3	13107	0
		1312	
be	1/4	26214	0
		528	
l2	1/4	26214	0
		1312	
af	1/4	26214	0
		3936	
l1	1/4	13107	0
		1312	
h2	1/4	26214	0
		3936	
ef	1/4	26214	0
		3936	
h1	1/4	13107	0
		1312	
nc	1/4	13107	0
		1312	
be	1/5	26214	0
		528	
l2	1/5	26214	0
		1312	
af	1/5	26214	0
		3936	
l1	1/5	13107	0

		1312	
h2	1/5	26214	0
		3936	
ef	1/5	26214	0
		3936	
h1	1/5	13107	0
		1312	
nc	1/5	13107	0
		1312	
be	1/6	26214	0
		528	
l2	1/6	26214	0
		1312	
af	1/6	26214	0
		3936	
l1	1/6	13107	0
		1312	
h2	1/6	26214	0
		3936	
ef	1/6	26214	0
		3936	
h1	1/6	13107	0
		1312	
nc	1/6	13107	0
		1312	
be	1/7	26214	0
		528	
l2	1/7	26214	0
		1312	
af	1/7	26214	0
		3936	
l1	1/7	13107	0
		1312	
h2	1/7	26214	0
		3936	
ef	1/7	26214	0
		3936	
h1	1/7	13107	0
		1312	
nc	1/7	13107	0
		1312	
be	1/8	26214	0
		528	
l2	1/8	26214	0
		1312	
af	1/8	26214	0
		3936	
l1	1/8	13107	0
		1312	
h2	1/8	26214	0
		3936	
ef	1/8	26214	0
		3936	
h1	1/8	13107	0
		1312	
nc	1/8	13107	0
		1312	
be	1/9	26214	0
		528	
l2	1/9	26214	0
		1312	
af	1/9	26214	0
		3936	
l1	1/9	13107	0

		1312	
h2	1/9	26214	0
		3936	
ef	1/9	26214	0
		3936	
h1	1/9	13107	0
		1312	
nc	1/9	13107	0
		1312	
be	1/10	26214	0
		528	
l2	1/10	26214	0
		1312	
af	1/10	26214	0
		3936	
l1	1/10	13107	0
		1312	
h2	1/10	26214	0
		3936	
ef	1/10	26214	0
		3936	
h1	1/10	13107	0
		1312	
nc	1/10	13107	0
		1312	
be	1/11	26214	0
		528	
l2	1/11	26214	0
		1312	
af	1/11	26214	0
		3936	
l1	1/11	13107	0
		1312	
h2	1/11	26214	0
		3936	
ef	1/11	26214	0
		3936	
h1	1/11	13107	0
		1312	
nc	1/11	13107	0
		1312	
be	1/12	26214	0
		528	
l2	1/12	26214	0
		1312	
af	1/12	26214	0
		3936	
l1	1/12	13107	0
		1312	
h2	1/12	26214	0
		3936	
ef	1/12	26214	0
		3936	
h1	1/12	13107	0
		1312	
nc	1/12	13107	0
		1312	
be	1/X1	26214	0
		528	
l2	1/X1	26214	0
		1312	
af	1/X1	26214	0
		3936	
l1	1/X1	13107	0

		1312	
h2	1/X1	26214	0
		3936	
ef	1/X1	26214	0
		3936	
h1	1/X1	13107	0
		1312	
nc	1/X1	13107	0
		1312	
be	1/X2	26214	0
		528	
l2	1/X2	26214	0
		1312	
af	1/X2	26214	0
		3936	
l1	1/X2	13107	0
		1312	
h2	1/X2	26214	0
		3936	
ef	1/X2	26214	0
		3936	
h1	1/X2	13107	0
		1312	
nc	1/X2	13107	0
		1312	
be	1/X3	26214	0
		528	
l2	1/X3	26214	0
		1312	
af	1/X3	26214	0
		3936	
l1	1/X3	13107	0
		1312	
h2	1/X3	26214	0
		3936	
ef	1/X3	26214	0
		3936	
h1	1/X3	13107	0
		1312	
nc	1/X3	13107	0
		1312	
be	1/X4	26214	0
		528	
l2	1/X4	26214	0
		1312	
af	1/X4	26214	0
		3936	
l1	1/X4	13107	0
		1312	
h2	1/X4	26214	0
		3936	
ef	1/X4	26214	0
		3936	
h1	1/X4	13107	0
		1312	
nc	1/X4	13107	0
		1312	
be	MCast	26214	0
		528	
l2	MCast	26214	0
		528	
af	MCast	26214	0
		528	
l1	MCast	13107	0

h2	MCast	528	
		26214	0
ef	MCast	528	
		26214	0
h1	MCast	528	
		13107	0
nc	MCast	528	
		13107	0
		528	
=====			
IPV4 GRE Fragment Reassembly Queue Groups			
=====			
FC-Maps ID	Profile ID	0.MBS(buf)	Depth(buf)
		0.CBS(buf)	

be	1	26214	0
		528	
l2	1	26214	0
		1312	
af	1	26214	0
		3936	
l1	1	13107	0
		1312	
h2	1	26214	0
		3936	
ef	1	26214	0
		3936	
h1	1	13107	0
		1312	
nc	1	13107	0
		1312	
....			
be	2	6214	0
		528	
l2	2	6214	0
		312	
....			
=====			
*A:ALU-1#			

Output example (2-port 10GigE (Ethernet) Adapter card)

The following output is an example of showing both v-port and ring buffer pool information for a 2-port 10GigE (Ethernet) Adapter card. [Table 48: Buffer pool field descriptions](#) describes the fields.

=====			
V-Port Buffer Pool Information			
=====			
Pool Total	: 524287 buffers	Buffer Size	: 512 bytes
Pool Shared	: 206119 buffers	Pool Resv	: 318168 buffers
Pool Total In Use	: 146 buffers		
Pool Exhaustion Drop	: 0		
=====			
Ring Buffer Pool Information			
=====			
Pool Total	: 262143 buffers	Buffer Size	: 768 bytes
Pool Shared	: 233663 buffers	Pool Resv	: 28480 buffers
Pool Total In Use	: 0 buffers		
Pool Exhaustion Drop	: 0		

```
=====
Network Ingress Queues
=====
```

```
-----
FC-Maps          Dest      0.MBS(buf) Depth(buf)
                   0.CBS(buf)
-----
```

```
No Match Found
=====
```

```
=====
Network Egress Queues
=====
```

```
-----
FC-Maps ID          0.MBS(buf) Depth(buf)
                   0.CBS(buf)
-----
```

```
be      1/3/v-port      26214      0
                   528
l2      1/3/v-port      26214      0
                   1312
af      1/3/v-port      26214      0
                   3936
l1      1/3/v-port      13107      0
                   1312
h2      1/3/v-port      26214      0
                   3936
ef      1/3/v-port      26214      0
                   3936
h1      1/3/v-port      13107      0
                   1312
nc      1/3/v-port      13107      0
                   1312
=====
```

```
=====
Network Ring Add-drop Port Queues
=====
```

```
-----
Queue ID          0.MBS(buf) Depth(buf)
                   0.CBS(buf)
-----
```

```
1      13107      0
                   262
2      13107      0
                   655
3      13107      0
                   1966
4      6553      0
                   655
5      13107      0
                   1966
6      13107      0
                   1966
7      6553      0
                   655
8      6553      0
                   655
=====
```

```
=====
Network Ring Port Egress Queues
=====
```

```
-----
Queue ID          0.MBS(buf) Depth(buf)
-----
```

		0.CBS(buf)	

1	1/3/1	13107 262	0
2	1/3/1	13107 655	0
3	1/3/1	13107 1966	0
4	1/3/1	6553 655	0
5	1/3/1	13107 1966	0
6	1/3/1	13107 1966	0
7	1/3/1	6553 655	0
8	1/3/1	6553 655	0
1	1/3/2	13107 262	0
2	1/3/2	13107 655	0
3	1/3/2	13107 1966	0
4	1/3/2	6553 655	0
5	1/3/2	13107 1966	0
6	1/3/2	13107 1966	0
7	1/3/2	6553 655	0
8	1/3/2	6553 655	0
=====			

Table 48: Buffer pool field descriptions

Label	Description
Buffer Pool Information (not applicable for the 10-port 1GigE/1-port 10GigE X-Adapter card)	
Ingress Buffer Pool Information (applies only to the 10-port 1GigE/1-port 10GigE X-Adapter card)	
Egress Buffer Pool Information (applies only to the 10-port 1GigE/1-port 10GigE X-Adapter card)	
V-Port Buffer Pool Information (applies only to the 2-port 10GigE (Ethernet) Adapter card and 2-port 10GigE (Ethernet) module)	
Ring Buffer Pool Information (applies only to the 2-port 10GigE (Ethernet) Adapter card and 2-port 10GigE (Ethernet) module)	
Pool Total	The total number of available buffers
Pool Shared	The number of buffers that can be shared

Label	Description
Pool Total In Use	The total number of buffers in use, in real time
Pool Exhaustion Drop	The number of packets dropped due to pool exhaustion
Buffer Size	The buffer size supported by the adapter card: 512 or 2304 bytes
Pool Resv	The number of packets reserved or committed
Access Ingress Queues	
Access Egress Queues	
Name	For access ingress queues associated with $N > 1$ Apipes, the format of the queue identifier is: Service-Id->Sap Aggregation Group Name->Queue Number All other queues use the following: Service-Id->Sap-Id->Queue Number
FC-Maps	The forwarding class-to-queue mappings
O.MBS (buf)	The maximum operational buffers in the queue
O.CBS (buf)	The committed operational buffers in the queue
Depth (buf)	The queue occupancy (the number of buffer units currently queued), in real time For byte-based WRED, the number of Depth buffers may exceed the number of MBS buffers
Network Ingress Queues	
FC-Maps	The forwarding class-to-queue mappings
Dest	The destination MDA identifier
O.MBS (buf)	The maximum operational buffers in the queue
O.CBS (buf)	The committed operational buffers in the queue
Depth (buf)	The queue occupancy (the number of buffer units currently queued), in real time For byte-based WRED, the number of Depth buffers may exceed the number of MBS buffers
Network Egress Queues	
FC-Maps ID	The forwarding class identifier

Label	Description
	Entries with the <i>port-id:tag->int-name</i> format indicate the use of per-VLAN shapers
O.MBS (buf)	The maximum operational buffers in the queue
O.CBS (buf)	The committed operational buffers in the queue
Depth (buf)	The queue occupancy (the number of buffer units currently queued), in real time For byte-based WRED, the number of Depth buffers may exceed the number of MBS buffers
IPv4 GRE Fragment Reassembly Queue Groups (not applicable for the 2-port 10GigE (Ethernet) Adapter card)	
FC-Maps ID	The forwarding class identifier
Profile ID	The IP reassembly profile identifier
O.MBS (buf)	The maximum operational buffers in the queue
O.CBS (buf)	The committed operational buffers in the queue
Depth (buf)	The queue occupancy (the number of buffer units currently queued), in real time

7 Slope QoS policies

This chapter provides information to configure slope QoS policies using the command line interface.

Topics in this chapter include:

- [Overview](#)
- [Basic configuration](#)
- [Service management tasks](#)
- [Slope QoS policy command reference](#)

7.1 Overview

Random early detection (RED) and weighted random early detection (WRED) queue management policies are associated with queues and can be created at both access and network ports and in both directions (that is, ingress and egress). The main difference is that with WRED, there can be more than one slope curve managing the fill rate of the same queue. One curve manages the discards on high-priority traffic, and another curve manages the discards on low-priority traffic. For more information, see the [Slope policies \(WRED and RED\)](#).

On all adapter cards and platforms except the Gen-3 Ethernet adapter cards and platforms (such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X with Ethernet ports), random discards (WRED) are buffer-based (that is, buffer-count in the queue is used to calculate the discard threshold). On the Gen-3 Ethernet adapter cards and platforms, random discards are byte-based (that is, payload-count (bytes) in the queue is used to calculate the discard threshold). The 7705 SAR-X with a TDM MDA uses buffer-based WRED.

7.2 Basic configuration

This section contains the following topics related to creating and applying slope QoS policies:

- [Creating a slope QoS policy](#)
- [Applying slope policies](#)
- [Default slope policy values](#)

A basic slope QoS policy must conform to the following rules:

- Each slope policy must have a unique policy ID.
- High slope and low slope are shut down (default).
- Default values can be modified but parameters cannot be deleted.

7.2.1 Creating a slope QoS policy

Configuring and applying QoS policies is optional. If no QoS policy is explicitly defined, a default QoS policy is applied.

To create a new slope policy, you must define the following:

- a slope policy name – the system does not dynamically assign a name
- a description – a brief description of the of policy
- the high slope for the high-priority WRED/RED slope graph
- the low slope for the low-priority WRED/RED slope graph

Use the following CLI syntax to configure a slope policy:

CLI syntax:

```
config>qos#
  slope-policy name
    description description-string
    high-slope
      max-avg percent
      max-prob percent
      start-avg percent
      no shutdown
    low-slope
      max-avg percent
      max-prob percent
      start-avg percent
      no shutdown
```

Example:

```
*A:ALU-1#
configure qos slope-policy "SlopePolicy1" create
config>qos>slope-policy$ description "Test1"
config>qos>slope-policy$ high-slope
config>qos>slope-policy>high-slope$ max-avg 90
config>qos>slope-policy>high-slope$ max-prob 60
config>qos>slope-policy>high-slope$ start-avg 90
config>qos>slope-policy>high-slope$ shutdown
config>qos>slope-policy>high-slope$ exit
config>qos>slope-policy$ low-slope
config>qos>slope-policy>low-slope$ max-avg 75
config>qos>slope-policy>low-slope$ max-prob 40
config>qos>slope-policy>low-slope$ start-avg 75
config>qos>slope-policy>low-slope$ exit
config>qos>slope-policy$ exit
*A:ALU-1#
```

The following output shows the configuration for SlopePolicy1:

```
*A:ALU-1>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
      slope-policy "SlopePolicy1" create
        description "Test1"
        high-slope
          shutdown
          start-avg 90
```

```
        max-prob 60
    exit
    low-slope
        shutdown
        start-avg 75
        max-prob 40
    exit
exit
#-----
```

7.2.2 Applying slope policies

Slope policies are applied to network and access egress and ingress queues.
Use the following CLI syntax:

CLI syntax:

```
config>qos>network-queue>queue>slope-policy name
config>qos>sap-ingress>queue>slope-policy name
config>qos>sap-egress>queue>slope-policy name
```

7.2.3 Default slope policy values

The default slope policies are identified as default. The default policies cannot be edited or deleted. The following table displays the default slope policy parameters.

Table 49: Slope policy defaults

Field	Default
description	"Default slope policy"
high-slope	
shutdown	shutdown
start-avg	70
max-avg	90
max-prob	80
low-slope	
shutdown	shutdown
start-avg	50
max-avg	75
max-prob	80

The following output displays the default configuration:

```
A*A:ALU-1>config>qos# info detail
#-----
echo "QoS Policy Configuration"
#-----
...
    slope-policy "default" create
        description "Default slope policy."
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 80
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 80
        exit
    ...
```

7.3 Service management tasks

This section describes the following service management tasks:

- [Deleting QoS policies](#)
- [Copying and overwriting QoS policies](#)
- [Editing QoS policies](#)

7.3.1 Deleting QoS policies

A QoS policy cannot be deleted until it is removed from a network or access egress/ingress queue. Use the following CLI syntax:

CLI syntax:

```
config>qos>network-queue>queue>no slope-policy
config>qos>sap-ingress>queue>no slope-policy
config>qos>sap-egress>queue>no slope-policy
```

7.3.1.1 Removing a policy from the QoS configuration

Use the following CLI syntax to delete a slope policy:

CLI syntax:

```
config>qos# no slope-policy name
```

Example:

```
config>qos# no slope-policy SlopePolicy1
```

7.3.2 Copying and overwriting QoS policies

You can copy an existing slope policy, rename it with a new policy ID value, or overwrite an existing policy ID. The **overwrite** option must be specified or an error occurs if the destination policy ID exists.

Use the following syntax to overwrite an existing QoS slope policy.

CLI syntax:

```
config>qos# copy slope-policy source-policy-name dest-policy-name
[overwrite]
```

Example:

```
*A:ALU-1>config>qos# copy slope-policy SlopePolicy1 SlopePolicy2 overwrite
config>qos# exit
*A:ALU-2#
```

The following output displays the copied policies:

```
*A:ALU-2>config>qos# info detail
#-----
echo "QoS Policy Configuration"
#-----
...
    slope-policy "default" create
        description "Default slope policy."
        high-slope
            shutdown
            start-avg 70
            max-avg 90
            max-prob 80
        exit
        low-slope
            shutdown
            start-avg 50
            max-avg 75
            max-prob 80
        exit
    exit
    slope-policy "SlopePolicy2" create
        description "Test2"
        high-slope
            shutdown
            max-avg 100
            start-avg 100
            max-prob 75
        exit
        low-slope
            shutdown
            start-avg 75
            max-avg 75
            max-prob 40
        exit
    exit
    slope-policy "SlopePolicy1" create
        description "Test1"
        high-slope
            shutdown
            start-avg 90
            max-avg 60
            max-prob 90
```

```
        exit
        low-slope
            shutdown
            start-avg 75
            max-avg 75
            max-prob 40
        exit
    exit
    slope-policy "SlopePolicy2" create
        description "Test1"
        high-slope
            shutdown
            start-avg 90
            max-avg 60
            max-prob 90
        exit
        low-slope
            shutdown
            start-avg 75
            max-avg 75
            max-prob 40
        exit
    exit
    ...
```

7.3.3 Editing QoS policies

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

7.4 Slope QoS policy command reference

7.4.1 Command hierarchies

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

7.4.1.1 Configuration commands

```
config
- qos
  - [no] slope-policy name [create]
    - description description-string
    - no description
    - [no] high-slope
      - max-avg percent
      - no max-avg
      - max-prob percent
      - no max-prob
      - start-avg percent
      - no start-avg
      - [no] shutdown
    - [no] low-slope
      - max-avg percent
      - no max-avg
      - max-prob percent
      - no max-prob
      - start-avg percent
      - no start-avg
      - [no] shutdown
```

7.4.1.2 Operational commands

```
config
- qos
  - copy slope-policy src-name dst-name [overwrite]
```

7.4.1.3 Show commands

```
show
- qos
  - slope-policy [slope-policy-name] [detail]
```

7.4.2 Command descriptions

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

7.4.2.1 Configuration commands

- [Generic commands](#)
- [Slope policy QoS commands](#)
- [WRED/RED slope commands](#)

7.4.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>slope-policy

Description

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

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

Default

n/a

Parameters

description-string

a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

7.4.2.1.2 Slope policy QoS commands

slope-policy

Syntax

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

Context

config>qos

Description

This command enables the context to configure a QoS slope policy.

Default

slope-policy "default"

Parameters

name

the name of the slope policy

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

create

keyword used to create a slope policy

high-slope

Syntax

[no] high-slope

Context

config>qos>slope-policy

Description

This command enables the context to define the high-priority Weighted Random Early Detection (WRED) or Random Early Detection (RED) slope graph. Each queue supports a high-priority WRED/RED slope for managing access to the queue for high-priority or in-profile packets.

The **high-slope** parameters can be changed at any time and the affected queue high-priority WRED/RED slopes are adjusted appropriately.

The **no** form of this command resets the parameters under the **high-slope** command context to the default values. If the parameters are set to the default values, the **high-slope** mode will not appear in **save config** and **show config** outputs unless the **detail** parameter is used.

low-slope

Syntax

[no] low-slope

Context

config>qos>slope-policy

Description

This command enables the context to define the low-priority WRED/RED slope graph. Each queue supports a low-priority WRED/RED slope for managing access to the queue for low-priority or out-of-profile packets.

The **low-slope** parameters can be changed at any time and the affected queue low-priority WRED/RED slopes are adjusted appropriately.

The **no** form of this command resets the parameters under the low-slope command context to the default values. If the parameters are set to the default values, the **low-slope** mode will not appear in **save config** and **show config** outputs unless the **detail** parameter is used.

7.4.2.1.3 WRED/RED slope commands

max-avg

Syntax

max-avg percent

no max-avg

Context

config>qos>slope-policy>high-slope

config>qos>slope-policy>low-slope

Description

This command sets the low-priority or high-priority WRED or RED slope position for the queue average usage value where the packet discard probability rises directly to one (maxThreshold or maxT). The *percent* parameter is expressed as a percentage of the maximum queue depth (buffer-based WRED) or a percentage of maximum payload of the queue (payload-based WRED).

The **no** form of this command restores the **max-avg** value to the default setting. If the current **start-avg** setting is larger than the default, an error will occur and the **max-avg** setting will not be changed to the default.

Default

max-avg 90 (the high slope default is 90% queue usage before discard probability is 1)

max-avg 75 (the low slope default is 75% queue usage before discard probability is 1)

Parameters

percent

the percentage of the maximum queue depth at which point the drop probability becomes 1. The value entered must be greater than or equal to the current setting of **start-avg**. If the entered value is smaller than the current value of **start-avg**, an error will occur and no change will take place.

Values 0 to 100

max-prob

Syntax

max-prob *percent*

no max-prob

Context

config>qos>slope-policy>high-slope

config>qos>slope-policy>low-slope

Description

This command sets the high-priority or low-priority WRED/RED maximum discard probability (maxDP) at slope position **max-avg**. The *percent* parameter is expressed as a percentage of packet discard probability where always discard is a probability of 1. A **max-prob** value of 80 represents 80% of 1, or a packet discard probability of 0.8.

For bridging domain queues on the 2-port 10GigE (Ethernet) Adapter card and 2-port 10GigE (Ethernet) module, the queues support the following DP (discard probability) values: 0%, 1%, 2%, 3%, 4%, 5%, 6%, 7%, 8%, 9%, 10%, 25%, 50%, 75%, and 100%. User-configured values are rounded down to match these DP values.

For example, configuring a DP to be 74% means that the actual value used is 50%.

The **no** form of this command restores the **max-prob** value to the default setting.

Default

max-prob 80 (80% maximum drop probability corresponding to the **max-avg**)

Parameters

percent

the maximum drop probability percentage corresponding to the **max-avg**, expressed as a decimal integer

Values 0 to 100

shutdown

Syntax

[no] shutdown

Context

config>qos>slope-policy>high-slope

config>qos>slope-policy>low-slope

Description

This command enables or disables the administrative status of the WRED/RED slope.

By default, all slopes are shut down and have to be explicitly enabled (**no shutdown**).

The **no** form of this command administratively enables the WRED/RED slope.

Default

shutdown (the WRED/RED slope disabled, implying a zero (0) drop probability)

start-avg

Syntax

start-avg *percent*

no start-avg

Context

config>qos>slope-policy>high-slope

config>qos>slope-policy>low-slope

Description

This command sets the high-priority or low-priority WRED/RED slope position for the queue average usage value where the packet discard probability starts to increase above zero (minThreshold or minT). The *percent* parameter is expressed as a percentage of the maximum queue depth (buffer-based WRED) or a percentage of maximum payload of the queue (payload-based WRED).

The **no** form of this command restores the **start-avg** value to the default setting. If the **max-avg** setting is smaller than the default, an error will occur and the **start-avg** setting will not be changed to the default.

Parameters

<i>percent</i>	the percentage of the maximum queue depth where the packet discard probability starts to increase above zero
Values	0 to 100
Default	50

7.4.2.2 Operational commands

copy

Syntax

copy slope-policy *src-name dst-name* [overwrite]

Context

config>qos

Description


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

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

Parameters

slope-policy <i>src-name dst-name</i>	indicates that the source policy ID and the destination policy ID are slope policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.
overwrite	specifies that the existing destination policy is to be replaced. Everything in the existing destination policy will be overwritten with the contents of the source policy. If overwrite is not specified, an error will occur if the destination policy ID exists.

7.4.2.3 Show commands

 **Note:** The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

slope-policy

Syntax

slope-policy [*slope-policy-name*] [**detail**]

Context

show>qos

Description

This command displays slope policy information.

Parameters

- slope-policy-name*
the name of the slope policy
- detail**
displays detailed information about the slope policy

Output

The following output is an example of slope policy information, and [Table 50: Slope policy field descriptions](#) describes the fields.

Output example

```
*A:ALU-1# show qos slope-policy SlopePolicy1 detail
=====
QoS Slope Policy
=====
Policy          : SlopePolicy1
Description     : Test1
Time Avg       : 3
-----
High Slope Parameters
-----
Start Avg      : 90
Max Avg       : 90
Admin State   : Disabled
Max Prob.    : 60
-----
Low Slope Parameters
-----
Start Avg      : 75
Max Avg       : 75
Admin State   : Disabled
Max Prob.    : 40
-----
Associations
-----
Object Type   Object Id   Queue
-----
sap-ingress   1         1
sap-ingress   8         1
sap-ingress   8         2
sap-egress    1         1
network-queue default    1
network-queue default    2
network-queue default    3
network-queue default    4
```

```

network-queue default 5
network-queue default 6
network-queue default 7
network-queue default 8
network-queue default 9
network-queue default 10
network-queue default 11
network-queue default 12
network-queue default 13
network-queue default 14
network-queue default 15
network-queue default 16
=====
*A:ALU-1#

```

Table 50: Slope policy field descriptions

Label	Description
Policy	The ID that uniquely identifies the policy
Description	A text string that helps identify the policy's context in the configuration file
Time Avg	The time average factor, which is the exponential weight factor used in calculating the average queue size. The <i>time_average_factor</i> parameter is non-user-configurable, and is set to a system-wide default value of 3.
High Slope Parameters	
Start Avg	The high-priority WRED/RED slope position for the queue average usage value where the packet discard probability starts to increase above zero
Max Avg	The high-priority WRED or RED slope position for the queue average usage value where the packet discard probability rises directly to one
Admin State	enabled - the administrative status of the WRED/RED slope is enabled
	disabled - the administrative status of the WRED/RED slope is disabled
Max Prob.	The high-priority WRED/RED maximum discard probability (at slope position max-avg)
Low Slope Parameters	
Start Avg	The low-priority WRED/RED slope position for the queue average usage value where the packet discard probability starts to increase above zero

Label	Description
Max Avg	The low-priority WRED or RED slope position for the queue average usage value where the packet discard probability rises directly to one
Admin State	enabled - the administrative status of the WRED/RED slope is enabled
	disabled - the administrative status of the WRED/RED slope is disabled
Max Prob.	The low-priority WRED/RED maximum discard probability (at slope position max-avg)
Associations	
Object Type	The type of object using the specified slope policy
Object Id	The identifier of the object using the specified slope policy
Queue	The number of the queue using the specified slope policy

8 ATM QoS traffic descriptor profiles

This chapter provides information to configure QoS traffic descriptor profiles using the command line interface.

Topics in this chapter include:

- [ATM traffic descriptor profiles](#)
- [Basic configuration](#)
- [Service management tasks](#)
- [ATM QoS policy command reference](#)

8.1 ATM traffic descriptor profiles

This section provides a description of support for ATM QoS policy features. Each traffic descriptor defines the expected rates and characteristics of traffic.

8.1.1 ATM traffic management

The 7705 SAR supports the ATM Forum Traffic Management Specification Version 4.1.

This section contains the following topics related to the QoS features for ATM Permanent Virtual Connections (PVCs):

- [ATM service categories](#)
- [ATM traffic descriptors and QoS parameters](#)
- [ATM policing](#)
- [Shaping](#)
- [ATM queuing and scheduling](#)
- [Congestion avoidance](#)

8.1.1.1 ATM service categories

The 7705 SAR supports the following service categories:

- CBR – constant bit rate
- rt-VBR – real-time variable bit rate
- nrt-VBR –non-real-time variable bit rate
- UBR/UBR+MIR – unspecified bit rate with minimum cell rate (UBR is a special case of UBR+MIR where MIR=0)

8.1.1.2 ATM traffic descriptors and QoS parameters

The following table shows the ATM traffic descriptors supported on the 7705 SAR.

Table 51: ATM traffic descriptors

Service category	Traffic descriptors
CBR	P0_1
	PIR in kb/s (applies to CLP=0 and CLP=1 flows)
rt-VBR and nrt-VBR	P0_1 and S0_1
	PIR in kb/s (applies to CLP=0 and CLP=1 flows)
	SIR in kb/s (applies to CLP=0 and CLP=1 flows)
	MBS in cells (applies to CLP=0 and CLP=1 flows)
	P0_1 and S0
	PIR in kb/s (applies to CLP=0 and CLP=1 flows; non-conforming CLP=0 cells are discarded)
	SIR in kb/s (applies to CLP=0 flow only)
	MBS in cells (applies to CLP=0 flow only)
	P0_1 and S0_Tag
	PIR in kb/s (applies to CLP=0 and CLP=1 flows; non-conforming CLP=0 flows are tagged to CLP=1 flows)
UBR/UBR+MIR	P0_1
	PIR in kb/s (applies to CLP=0 and CLP=1 flows)
	MIR in kb/s (applies to CLP=0 and CLP=1 flows)

8.1.1.3 ATM policing

The policing option, when enabled, applies only to ingress traffic. Similarly, the shaping option, if enabled, applies only to egress traffic. For example, if a traffic descriptor has both policing and shaping enabled, the policing option is enforced for the ingress traffic, while the shaping option is enforced for the egress traffic. The following ATM service category conformance definitions are supported:

- P0_1 – CBR
- P0_1 and S0_1 – VBR.1
- P0_1 and S0 – VBR.2
- P0_1 and S0_Tag – VBR.3

P represents the peak rate, S represents the sustained rate, 0 or 1 represents the CLP value to which policing is applied if the cell is non-conforming, and Tag indicates that the CLP value has changed from 0 to 1. For example:

- P0_1 – means that policing is applied to non-conforming peak rate cells with CLP values of 0 or 1
- P0_1 and S0_Tag – means that policing is applied to non-conforming peak rate cells with CLP values of 0 or 1 and policing is applied to non-conforming sustained rate cells with a CLP value of 0 with the action to change the CLP value to 1

8.1.1.4 Shaping

ATM layer egress shaping is supported for CBR, rt-VBR, and nrt-VBR VCs. A CBR VC is shaped to a single leaky bucket with the parameter PIR from the traffic descriptor. An rt-VBR VC or an nrt-VBR VC is shaped to two leaky buckets with parameters PIR and SIR, BT from the traffic descriptor, where BT is the burst tolerance and is a function of the MBS parameters configured by the user. The traffic rates and shaping parameter that are configured in the **sap-egress** QoS policy are not used at ATM SAPs.

To enforce the SLA, ATM layer ingress policing is supported at ingress, so no shaping is needed. At ingress, after optional policing is applied, packet level queue-based soft-policing is supported per the service ingress QoS policy applied to the ATM SAP.



Note: Shaping to the specified traffic descriptor in the ATM traffic descriptor profile is always enabled for CBR and rt-VBR VCs.

8.1.1.5 ATM queuing and scheduling

The 7705 SAR provides a per-VC queuing architecture on the 16-port T1/E1 ASAP Adapter card, 32-port T1/E1 ASAP Adapter card, and 2-port OC3/STM1 Channelized Adapter card with atm/ima encapsulation. The 7705 SAR provides a per-VC queuing architecture on the 4-port OC3/STM1 Clear Channel Adapter card and 4-port DS3/E3 Adapter card with atm encapsulation. In the egress direction toward the ATM port, the scheduling priority at the ATM layer is as follows:

- CBR VCs are always shaped and are scheduled with strict priority over all other service categories.
- rt-VBR VCs are always shaped and are scheduled next with strict priority over nrt-VBR and UBR VCs.
- nrt-VBR shaped VCs are scheduled next with strict priority over nrt-VBR unshaped VCs and UBR VCs.
- nrt-VBR unshaped VCs and UBR VCs are scheduled as a common class. Scheduling among these VCs is done using a weighted round robin (WRR) scheduler, where the weight of each VC is determined by the configured SIR for nrt-VBR and by the MIR for UBR VCs. The scheduling is work-conserving, so each VC has access to excess bandwidth in proportion to its SIR/MIR. Under congestion, the performance of each VC degrades proportionally to the weight of the VC.

8.1.1.6 Congestion avoidance

Congestion and potential discards are performed on a per-forwarding class basis in the SAP queues in the CSM.

8.2 Basic configuration

This section contains the following topics related to creating and applying ATM QoS policies:

- [Creating an ATM traffic descriptor profile QoS policy](#)
- [Applying ATM traffic descriptor profile policies](#)
- [Default ATM traffic descriptor profile policy values](#)

A basic ATM QoS traffic descriptor profile must conform to the following rules:

- Each policy must have a unique policy ID.
- Default values can be modified but parameters cannot be deleted.

8.2.1 Creating an ATM traffic descriptor profile QoS policy

Configuring and applying QoS policies and profiles other than the default policy is optional. To create an ATM QoS traffic descriptor profile, perform the following:

- assign a policy ID (policy number) – the system does not dynamically assign an ID
- include a description – provides a brief overview of policy features
- configure traffic attributes of the ATM traffic profile
- determine whether egress shaping should occur

Use the following CLI syntax to configure an **atm-td-profile** policy.

CLI syntax:

```
config>qos#
  atm-td-profile traffic-desc-profile-id
    description description-string
    descriptor-type {P0_1|P0_1andS0_Tag|P0_1andS0|P0_1andS0_1}
    [no] policing
    service-category service-category
    [no] shaping
    traffic [sir sir-val] [pir pir-val] [mir mir-val] [mbs mbs-val]
    [cdvt cdvt-val]
```

The following output displays an example of the command syntax.

Example:

```
*A:ALU-1# config qos
config>qos# atm-td-profile 3 create
config>qos>atm-td-profile$ description "ATM TD profile3"
config>qos>atm-td-profile$ service-category rt-vbr
config>qos>atm-td-profile$ descriptor-type P0_1andS0_1
config>qos>atm-td-profile$ policing
config>qos>atm-td-profile$ shaping
config>qos>atm-td-profile$ traffic sir 500 pir 500 mbs 500 cdvt 500
config>qos>atm-td-profile$ exit
config>qos# exit
```

The following output displays the profile configuration for ATM TD profile 3:

```
*A:ALU-1>config>qos# info
```

```
#-----
echo "QoS Policy Configuration"
#-----
...
    atm-td-profile 3 create
        description "ATM TD profile3"
        service-category rt-vbr
        traffic sir 500 pir 500 mbs 500 cdvt 500
        policing
    exit
...
*A:ALU-1
```

8.2.2 Applying ATM traffic descriptor profile policies

ATM QoS traffic descriptor profiles are applied to ATM VLL (Apipe) SAPs. ATM QoS traffic descriptor profiles can also be applied on egress to ATM VCC SAPs that are members of a SAP aggregation group. You cannot apply ATM QoS traffic descriptor profiles on ingress to a SAP in a SAP aggregation group; the profile is set to the default (1).

8.2.2.1 ATM VLL (Apipe) SAPs

Use the following CLI syntax to apply ATM QoS traffic descriptor profile policies to Apipe SAPs on ingress and egress.

CLI syntax:

```
config>service>apip>sap# atm
    egress
        traffic-desc traffic-desc-profile-id
    ingress
        traffic-desc traffic-desc-profile-id
```

8.2.3 Default ATM traffic descriptor profile policy values

The default ATM QoS traffic descriptor profile is 1. The default profile cannot be edited or deleted.

The following table shows the ATM TD profile defaults.

Table 52: ATM TD profile defaults

Field	Default
atm-td-profile traffic-desc-profile-id	1
description	Default Traffic Descriptor
descriptor-type	Based on service category: CBR: P0_1 UBR: P0_1 UBR+MIR: P0_1

Field	Default
	rt-VBR or nrt-VBR: P0_1 and S0_1
policing	No policing
service-category	UBR
traffic	No traffic
shaping	No shaping

The following output displays the default configuration:

```
A:ALU-1>config>qos# info detail
-----
Echo "QoS Policy Configuration"
-----
    atm-td-profile 1 create
        description "Default Traffic Descriptor"
        service-category ubr
        no traffic
        no policing
        descriptor-type P0_1
        no shaping
    exit
...
```

8.3 Service management tasks

This section describes the following ATM Traffic Descriptor Profile service management tasks:

- [Removing an ATM traffic descriptor profile from the QoS configuration](#)
- [Copying and overwriting an ATM traffic descriptor profile](#)
- [Editing QoS policies](#)

8.3.1 Removing an ATM traffic descriptor profile from the QoS configuration

The default ATM traffic descriptor profile cannot be deleted.

To delete an ATM QoS traffic descriptor profile, enter the following command:

CLI syntax:

```
config>qos# no atm-td-profile traffic-desc-profile-id
```

Example:

```
config>qos# no atm-td-profile 3
```

8.3.2 Copying and overwriting an ATM traffic descriptor profile

You can copy an existing profile, rename it with a new profile ID value, or overwrite an existing profile ID. The **overwrite** option must be specified or an error occurs if the destination profile ID exists.

CLI syntax:

```
config>qos# copy atm-td-profile src-prof dst-prof [overwrite]
```

Example:

```
*A:ALU-1#>config>qos# copy atm-td-profile 2 3
A:ALU-48>config>qos# copy atm-td-profile 2 3 overwrite
A:ALU-48>config>qos#
```

8.3.3 Editing QoS policies

You can change existing policies and entries in the CLI. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors, copy the policy to a work area, make the edits, and then write over the original policy.

8.4 ATM QoS policy command reference

8.4.1 Command hierarchies

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

8.4.1.1 Configuration commands

```
config
- qos
  - [no] atm-td-profile traffic-desc-profile-id [create]
    - description description-string
    - no description
    - descriptor-type type
    - [no] policing
    - service-category service-category
    - [no] shaping
    - traffic [sir sir-val] [pir pir-val] [mir mir-val] [mbs mbs-val] [cdvt cdvt-val]
    - no traffic
```

8.4.1.2 Operational commands

```
config
- qos
  - copy atm-td-profile src-prof dst-prof [overwrite]
```

8.4.1.3 Show commands

```
show
- qos
  - atm-td-profile [traffic-desc-profile-id] [detail]
- service
  - sap-using [ingress | egress] atm-td-profile td-profile-id
  - sap-using [ingress | egress] qos-policy qos-policy-id
```


8.4.2 Command descriptions

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

8.4.2.1 Configuration commands

- [Generic commands](#)
- [ATM QoS policy commands](#)

8.4.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>atm-td-profile

Description

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

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

Default

n/a

Parameters

description-string

a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

8.4.2.1.2 ATM QoS policy commands

atm-td-profile

Syntax

[no] **atm-td-profile** *traffic-desc-profile-id* [create]

Context

config>qos

Description

This command is used to configure an ATM traffic descriptor profile as per ATM Forum Traffic Management Specification Version 4.1.

Traffic descriptor profiles are used to:

- define traffic management capabilities for ATM PVCCs
- calculate the total bandwidth consumed on a given port by all ATM PVCCs. The bandwidth taken by a PVCC is equal to:
- PIR for CBR PVCCs
- SIR for rt-vbr and nrt-vbr PVCCs
- MIR for UBR PVCC
- define ATM-level scheduling

The default traffic descriptor is preconfigured and non-modifiable. It cannot be deleted. All other traffic descriptor profiles must be explicitly created before use. The create keyword must follow each new profile configuration.

Any changes made to the existing profile, using any of the sub-commands, are applied immediately to all objects where this profile is applied (a small traffic interruption in data traffic will occur during the data plane reprogramming with the newly modified profile).

When many changes are required on a profile, it is recommended that the profile be copied to a work area profile ID. That work-in-progress profile can be modified until complete and then written over the original profile-id. Use the config>qos>copy command to maintain profiles in this manner.

The weight assigned to each non-shaped PVCC in the Deficit Round Robin Scheduler depends on the service category and traffic rates (see the [traffic](#) command for more details).

The no form of the command deletes a given traffic profile. The profile to be deleted must not be associated with any object (for example a SAP). If this condition is not met, the command will return an error.

Default

1 – the default traffic descriptor (UBR, no traffic, no shaping)

Parameters

- traffic-desc-profile-id

the index identifier for a traffic descriptor profile

Values1 to 1000
- create

keyword used to create an ATM traffic descriptor profile

descriptor-type

Syntax

descriptor-type type


Context

config>qos>atm-td-profile

Description

This command is used to specify the type of ATM traffic descriptor profile as per ATM Forum Traffic Management Specification Version 4.1.

[Table 53: Service category descriptor type default values](#) defines the **descriptor-type** default values based on service category, and [Table 54: Traffic descriptor type parameters](#) defines how the traffic parameters that are specified for the descriptor profiles are interpreted.



Note: Setting **descriptor-type** to a value that is not compatible with the service category as defined in the table results in an error message.

Table 53: Service category descriptor type default values

Service category	Default descriptor type
CBR	P0_1
UBR	P0_1
UBR with MIR	P0_1
rt-VBR or nrt-VBR	P0_1and S0_1

Table 54: Traffic descriptor type parameters

Descriptor type value	Rates interpretation	Applicable service categories
P0_1	PIR in kb/s; applies to CLP=0 and CLP=1 cell flows MIR in kb/s; applies to CLP=0 and CLP=1 cell flows	CBR, UBR, UBR with MIR

Descriptor type value	Rates interpretation	Applicable service categories
P0_1andS0	PIR in kb/s; applies to CLP=0 and CLP=1 cell flows; non-conforming CLP=0 cell flows are discarded SIR in kb/s; applies to CLP=0 cell flows only	rt-VBR and nrt-VBR
P0_1andS0_Tag	PIR in kb/s; applies to CLP=0 and CLP=1 cell flows; non-conforming CLP=0 cell flows are tagged to CLP=1 cell flows SIR in kb/s; applies to CLP=0 cell flows only	rt-VBR and nrt-VBR
P0_1andS0_1	PIR in kb/s; applies to CLP=0 and CLP=1 cell flows SIR in kb/s; applies to CLP=0 cell flows only	rt-VBR and nrt-VBR

Default

see [Table 53: Service category descriptor type default values](#)

Parameters

type

the ATM traffic descriptor profile type

Values P0_1, P0_1andS0_Tag, P0_1andS0, P0_1andS0_1

policing**Syntax**

[no] policing

Context

config>qos>atm-td-profile

Description

This command determines whether ingress traffic is policed. Policing is valid for CBR, rt-VBR, nrt-VBR, and UBR. This policing is cell-based.

Default

disabled

service-category

Syntax

service-category *service-category*

Context

config>qos>atm-td-profile

Description

This command is used to configure an ATM service category attribute of an ATM traffic descriptor profile.

Default

ubr

Parameters

service-category

The following table describes the supported ATM service categories on the 16-port T1/E1 ASAP Adapter card, 32-port T1/E1 ASAP Adapter card, 4-port DS3/E3 Adapter card, 4-port OC3/STM1 Clear Channel Adapter card, and 2-port OC3/STM1 Channelized Adapter card.

Table 55: ATM service categories

Service category	Description
CBR	Constant bit rate
rt-VBR	Real-time variable bit rate
nrt-VBR	Non-real-time variable bit rate
UBR	Unspecified bit rate without minimum desired dell rate (defined by specifying service category to be ubr, and MIR of 0)
UBR (with MIR)	Unspecified bit rate with non-zero minimum desired cell rate (defined by specifying service category to be ubr, and MIR > 0)

Changing the service category of a profile will reset all traffic attributes to their defaults (see the [traffic](#) command) and will cause reprogramming of the data path (with a small impact on user traffic) and a reset of VC statistics for all VCs using this traffic descriptor profile.

shaping

Syntax

[no] shaping

Context

config>qos>atm-td-profile

Description

This command determines whether egress shaping should occur. Shaping is only applied in the egress direction.

Default

The default is determined by the service category. The following table describes the default shaping values.

Table 56: Default shaping values

Applicable service category	Default shaping value	Comments
UBR	Disabled	Shaping cannot be enabled
CBR	Enabled	Shaping cannot be disabled
rt-VBR	Enabled	Shaping cannot be disabled
nrt-VBR	Enabled	Shaping cannot be disabled

traffic

Syntax

traffic [sir sir-val [pir pir-val] [mir mir-val] [mbs mbs-val] [cdvt cdvt-val]

no traffic

Context

config>qos>atm-td-profile

Description

This command is used to configure traffic attributes of an ATM traffic profile as per ATM Forum Traffic Management Specification Version 4.1.

The traffic parameters of a traffic descriptor that are configurable depend on the service category of this traffic descriptor profile. (See the [service-category](#) command.)

The following table defines which traffic descriptor parameters are applicable for what service category and what the configuration rules are between the parameters. "Y" indicates that the parameter can be configured for a given service category and will be defaulted if not provided, and an "N/A" indicates that the parameter cannot be configured for a given service category (an error will be returned). If an applicable parameter is not specified, the current value will be preserved.

Table 57: Service category traffic descriptor parameters

Service category	SIR	PIR	MBS	MIR	CDVT
CBR	N/A	Y	N/A	N/A	Y
rt-VBR	Y	Y (must be \geq SIR)	Y	N/A	Y
nrt-VBR	Y	Y (must be \geq SIR)	Y	N/A	Y
UBR	N/A	Y	N/A	N/A	Y
UBR with MIR	N/A	Y (must be \geq MIR)	N/A	Y (non-zero MIR specified)	Y

When a traffic descriptor profile is used to define egress scheduling, the following describes how traffic rates are used to derive scheduling weight:

The scheduling weight applies only to unshaped nrt-VBR and UBR. The scheduling weight is a value from 1 to 8. The scheduling weight is determined by the SIR value for nrt-VBR, and by the MIR value for UBR. The conversion from SIR/MIR to weight is as follows:

- Rate < 64K weight = 1
- Rate < 128K weight = 2
- Rate < 256K weight = 3
- Rate < 512K weight = 4
- Rate < 1024K weight = 5
- Rate < 1536K weight = 6
- Rate < 1920K weight = 7

Everything above 1920K will be assigned a weight of 8.

Since the 7705 SAR operates in cells/second with one cell granularity, PIR and SIR values programmed need to be converted to cells/second. When converting values to be used for the scheduler, the result is rounded up to the next cell when required by conversion.

When any of SIR, PIR, or MIR is greater than the physical maximum port/channel capacity for a given PVCC, then the maximum physical port/channel capacity is used in bandwidth accumulation and when configuring the hardware for that PVCC.

Hardware-enforceable MBS is in the inclusive range of 3 to 256 000 cells on all 7705 SAR ATM-capable adapter cards. Assigning an **atm-td-profile** with an MBS value outside of this range to a SAP is blocked. As well, once an **atm-td-profile** is assigned to a SAP, the CLI will block any change of the MBS value to a value outside of this range.

The **no** form of the command restores traffic parameters to their defaults for a given service category.

Default

The following table shows the ATM traffic parameter defaults.

Table 58: ATM traffic parameter defaults

Service category	Traffic parameter defaults
CBR:	
PIR	0
CDVT	250
rt-VBR and nrt-VBR	
PIR	4294967295
SIR	0
MBS	32
CDVT	250
UBR (note by default UBR is without MIR)	
PIR	0
CDVT	250

Parameters

sir-val

the sustained information rate (including cell overhead) in kb/s

Values 0 to 4294967295

pir-val

the peak information rate (including cell overhead) in kb/s

Values 0 to 4294967295

mir-val

the minimum desired information rate (including cell overhead) in kb/s

Values 0 to 4294967295

mbs-val

the maximum burst size in cells

Values 0 to 4294967295 (this range applies if the **atm-td-profile** is not assigned to a SAP)
3 to 256000 (this range applies if the **atm-td-profile** is assigned to a SAP, as described above)

cdvt-val

the cell delay variation tolerance (CDVT) in microseconds

Default CBR/RT-VBR/NRT-VBR/UBR = 250**Values** 0 to 4294967295

8.4.2.2 Operational commands

copy

Syntax

copy atm-td-profile *src-prof dst-prof* [**overwrite**]

Context

config>qos

Description

This command copies the source ATM traffic descriptor profile into the destination ATM profile. If the destination profile was already defined, the keyword **overwrite** must be appended for the copy to complete.

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

Parameters

src-prof dst-prof

indicates that the source profile ID and the destination profile ID are **atm-td-profile** IDs. Specify the source ID that the copy command will copy and specify the destination ID to which the command will duplicate the profile.

Values 1 to 1000

overwrite

specifies that the existing destination profile is to be replaced. Everything in the existing destination profile will be overwritten with the contents of the source profile. If **overwrite** is not specified, an error will occur if the destination profile ID exists.

ALU-48>config>qos# copy atm-td-profile 2 10

MINOR: CLI destination (10) exists use {overwrite}.

ALU-48>config>qos# copy atm-td-profile 2 10 overwrite

ALU-48>config>qos#

8.4.2.3 Show commands



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

atm-td-profile

Syntax

atm-td-profile [*traffic-desc-profile-id*] [**detail**]

Context

show>qos

Description

This command displays ATM traffic descriptor profile information.

Parameters

traffic-desc-profile-id
displays the ATM traffic descriptor profile

Values 1 to 1000

detail
displays detailed policy information including policy associations

Output

The following output is an example of ATM traffic descriptor profile information, and [Table 59: ATM traffic descriptor profile field descriptions](#) describes the fields.

Output example

```
*A:ALU-1>show>qos# atm-td-profile 3 detail
=====
Traffic Descriptor Profile (3)
=====
-----
TDP-id Description
Service Cat SIR PIR MIR MBS CDVT
-----
3 ATM TD profile3
RT-VBR 500 500 - 500 500
-----
TDP details
-----
Shaping : enabled
Policing : enabled
Descriptor-Type : P0_landS0_1
-----
Entities using TDP-3
-----
```

```
=====
*A:ALU-1>show>qos#
```

Table 59: ATM traffic descriptor profile field descriptions

Label	Description
Maximum Supported Profiles	The maximum number of ATM traffic descriptor profiles that can be configured on this system
Currently Configured Profiles	The number of currently configured ATM traffic descriptor profiles on this system
TDP-Id	The ID that uniquely identifies the traffic descriptor policy
Description	A text string that helps identify the policy's context in the configuration file
Service Category	The ATM service category
SIR	The sustained cell rate in kb/s
PIR	The peak cell rate in kb/s
MIR	The minimum desired cell rate in kb/s
MBS	The maximum burst size in cells
CDVT	The cell delay variation tolerance in microseconds
Shaping	Whether shaping is enabled or disabled for the traffic descriptor profile
Policing	Whether policing is enabled or disabled for the traffic descriptor profile
Descriptor Type	The descriptor type for the ATM TD profile
Entities using TDP-ID	The number of entities using the ATM traffic descriptor
-	The parameter is not applicable for the configured service category

sap-using

Syntax

sap-using [ingress | egress] atm-td-profile *td-profile-id*

sap-using [ingress | egress] qos-policy *qos-policy-id*

Context

show>service

Description

This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters

ingress

specifies matching an ingress policy

egress

specifies matching an egress policy

qos-policy-id

identifies the ingress or egress QoS policy for which to display matching SAPs or SAP aggregation groups. SAP aggregation groups are supported only on ATM VLLs.

Values 1 to 65535, or *qos-policy-name* (up to 64 characters)

td-profile-id

displays SAPs using the specified traffic description

Output

The following output is an example of SAP information using the atm-td-profile parameter, and [Table 60: SAP field descriptions](#) describes the fields.

Output example

```
*A:ALU-1>show>service# sap-using egress atm-td-profile 3
=====
Service Access Point Using ATM Traffic Profile 3
=====
PortId                SvcId      Ing.  Ing.  Egr.  Egr.  Adm  Opr
                   QoS    Fltr  QoS   Fltr
-----
No Matching Entries
```

The following output shows an example of SAPs and SAP aggregation groups that use the same QoS policy. SAPs that are members of a SAP aggregation group are not listed by their port identifiers but are listed under their respective group names. [Table 60: SAP field descriptions](#) describes the fields.

Output example

```
*B:A:179_121# show service sap-using ingress qos-policy 1
=====
Service Access Points Using Ingress Qos Policy 1
=====
PortId                SvcId      Ing.  Ing.  Egr.  Egr.  Adm  Opr
                   QoS    Fltr  QoS   Fltr
-----
1/1/3.1:0/100         100        1    none   1    none   Up   Down
1/1/3.1:3/100         100        1    none   1    none   Up   Down
```

```

1/2/5:200          200      1      none    1      none    Up    Up
1/1/9.1:3/300      300      1      none    1      none    Up    Up
bundle-ima-1/1.1:6/601 2001    1      none    1      none    Up    Up
bundle-ima-1/1.1:6/602 2002    1      none    1      none    Up    Up
-----
Number of SAPs: 6
-----
=====
SAP Aggregation Groups Using Ingress Qos Policy 1
=====
Group Name          SvcId      Ing.   Ing.   Egr.   Egr.   Adm   Opr
                   QoS      Fltr   QoS    Fltr
-----
"Group Name 1"      2003      1      n/a    1      n/a    n/a   n/a
sap:1/1/9.1:3/301      atm          1524   1524   Up    Up
sap:1/1/9.1:3/302      atm          1524   1524   Up    Up
sap:1/1/9.1:3/303      atm          1524   1524   Up    Up
"Group Name 2"      2004      1      n/a    1      n/a    n/a   n/a
sap:1/1/9.1:3/304      atm          1524   1524   Up    Up
sap:1/1/9.1:3/305      atm          1524   1524   Up    Up
sap:1/1/9.1:3/306      atm          1524   1524   Up    Up
-----
Number of SAP Aggregation Groups:2
-----

```

Table 60: SAP field descriptions

Label	Description
Service Access Points Using ATM Traffic Profile	
Service Access Points Using Ingress Qos Policy	
PortId	The ID of the access port where the SAP is defined
SvcId	The service identifier
Ing.QoS	The SAP ingress QoS policy number specified on the ingress SAP
Ing.Fltr	The SAP ingress filter number specified on the ingress SAP
Egr.QoS	The SAP egress QoS policy number specified on the egress SAP
Egr.Fltr	The SAP egress filter number specified on the egress SAP
Adm	The administrative state of the SAP
Opr	The operational state of the SAP
Number of SAPs	The number of individual SAPs that use the specified QoS policy
SAP Aggregation Groups Using Ingress Qos Policy	
Group Name	The SAP aggregation group name identifier
SvcId	The service identifier for the group

Label	Description
Ing.QoS	The ingress QoS policy number specified on the ingress SAP aggregation group
Ing.Fltr	The SAP ingress filter number specified on the ingress SAP
Egr.QoS	The egress QoS policy number specified on the egress SAP aggregation group
Egr.Fltr	The SAP egress filter number specified on the egress SAP
Adm	The administrative state of the SAP in the SAP aggregation group
Opr	The operational state of the SAP in the SAP aggregation group
Number of SAP Aggregation Groups	The number of SAP aggregation groups that use the specified QoS policy

9 QoS fabric profiles

This chapter provides information to configure QoS fabric profiles using the command line interface.

- [Basic configuration](#)
- [Service management tasks](#)
- [QoS fabric profile command reference](#)

9.1 Basic configuration

This section contains the following topic related to creating and applying QoS fabric policies:

- [Creating a QoS fabric profile](#)
- [Applying a QoS fabric profile](#)
- [Default fabric profile values](#)

A QoS fabric profile must conform to the following rules:

- Each profile must be associated with a unique policy ID.
- Either aggregate mode or per-destination mode must be assigned.

9.1.1 Creating a QoS fabric profile

Creating a QoS fabric profile other than the default policy ("default") is optional. To create a QoS fabric profile, perform the following:

- assign a policy ID (policy number) – the system does not dynamically assign an ID
- include an optional description of the policy
- assign the mode, either aggregate or per-destination. If no mode is assigned, the default aggregate mode is used.
- configure the to-fabric shaper rate
- (optional) for hierarchical QoS (H-QoS), configure the unshaped SAP CIR rate (see [Configuring per-SAP aggregate shapers and an unshaped SAP aggregate shaper \(H-QoS\)](#))

Use the following CLI syntax to configure a QoS fabric profile:

CLI syntax:

```
config>qos#
  fabric-profile policy-id aggregate-mode create
    aggregate-rate aggregate-rate [unshaped-sap-cir cir-rate]
    description description-string
  fabric-profile policy-id destination-mode create
    description description-string
    dest-mds slot/mda | multipoint
    rate mda-rate
```

The following example shows the command syntax for creating and configuring a destination-mode QoS fabric profile with an mda-rate of 400 Mb/s for destination adapter cards 1/1 through 1/6.

Example:

```
*A:7705:Dut-C# configure qos
fabric-profile 4 destination-mode create
config>qos>fabric-profile$ description "Fabric profile QoS policy 4"
config>qos>fabric-profile$ dest-mda 1/1
config>qos>fabric-profile>dest-mda$ rate 400000
config>qos>fabric-profile>dest-mda$ exit
config>qos>fabric-profile$ dest-mda 1/2
config>qos>fabric-profile>dest-mda$ rate 400000
config>qos>fabric-profile>dest-mda$ exit
config>qos>fabric-profile$ dest-mda 1/3
config>qos>fabric-profile>dest-mda$ rate 400000
config>qos>fabric-profile>dest-mda$ exit
config>qos>fabric-profile$ dest-mda 1/4
config>qos>fabric-profile>dest-mda$ rate 400000
config>qos>fabric-profile>dest-mda$ exit
config>qos>fabric-profile$ dest-mda 1/5
config>qos>fabric-profile>dest-mda$ rate 400000
config>qos>fabric-profile>dest-mda$ exit
config>qos>fabric-profile$ dest-mda 1/6
config>qos>fabric-profile>dest-mda$ rate 400000
config>qos>fabric-profile>dest-mda$ exit
config>qos>fabric-profile$ dest-mda multipoint
config>qos>fabric-profile>dest-mda$ rate 400000
config>qos>fabric-profile>dest-mda$ exit
config>qos>fabric-profile$ exit
*A:7705:Dut-C#
```

The following output displays the profile configuration for fabric profile QoS policy 4.

```
*A:7705:Dut-C#>config>qos# info detail
#-----
echo "QoS Policy Configuration"
#-----
...
fabric-profile 4 destination-mode create
description "Fabric profile QoS policy 4"
dest-mda 1/1
rate 400000
exit
dest-mda 1/2
rate 400000
exit
dest-mda 1/3
rate 400000
exit
dest-mda 1/4
rate 400000
exit
dest-mda 1/5
rate 400000
exit
dest-mda 1/6
rate 400000
exit
dest-mda multipoint
rate 400000
exit
```


9.1.2 Applying a QoS fabric profile

Fabric profiles do not apply to the Auxiliary Alarm card or the 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M, or 7705 SAR-Wx.

Use the following CLI syntax to assign a fabric profile on an adapter card.

CLI syntax:

```
config>card>mda#
network
  ingress
    fabric-policy fabric-policy-id
    queue-policy name
access
  ingress
    fabric-policy fabric-policy-id
no shutdown
```

9.1.3 Default fabric profile values

The following table shows the fabric profile default values.

Table 61: Fabric profile defaults

Field	Default
Policy-id	1
Mode	Aggregate-mode
Aggregate rate	200000 kb/s
CIR rate	0 kb/s
MDA rate	200000 kb/s

The following output displays the default configuration:

```
A:ALU-1>config>qos# info detail
-----
Echo "QoS Policy Configuration"
-----
...
fabric-profile 1 aggregate-mode create
    description "Default Fabric Profile QoS policy."
    aggregate-rate 200000 unshaped-sap-cir 0
exit
...
-----
```

9.2 Service management tasks

This section describes the following fabric profile service management tasks:

- [Removing a fabric profile from the QoS configuration](#)
- [Copying and overwriting a fabric profile](#)
- [Editing QoS policies](#)

9.2.1 Removing a fabric profile from the QoS configuration

The default fabric profile cannot be deleted.

To delete a fabric profile, enter the following command:

CLI syntax:

```
config>qos# no fabric-profile policy-id
```

Example:

```
config>qos# no fabric-profile 3
```

9.2.2 Copying and overwriting a fabric profile

You can copy an existing profile, rename it with a new profile ID value, or overwrite an existing profile ID. The **overwrite** option must be specified or an error occurs if the destination profile ID exists.

CLI syntax:

```
config>qos# copy fabric-profile src-prof dst-prof [overwrite]
```

Example:

```
*A:ALU-1#>config>qos# copy fabric-profile 2 3  
config>qos# copy fabric-profile 2 3 overwrite  
config>qos#
```

9.2.3 Editing QoS policies

You can change existing policies and entries in the CLI. The changes are applied immediately to the specified adapter card where the policy is applied. To prevent configuration errors, copy the policy to a work area, make the edits, and then write over the original policy.

9.3 QoS fabric profile command reference

9.3.1 Command hierarchies

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

9.3.1.1 Configuration commands

The **config>qos>fabric-profile** command applies only to the 7705 SAR-8 Shelf V2 and 7705 SAR-18; it does not apply to the fixed platforms. The **config>system>qos** command applies only to the fixed platforms, with the exception of the 7705 SAR-X. The 7705 SAR-X does not support fabric aggregate rate configuration; its shaper rate is set permanently to 16 Gb/s.

```
config
- qos
  - fabric-profile policy-id [aggregate-mode | destination-mode] [create]
  - no fabric-profile policy-id
    - aggregate-rate aggregate-rate [unshaped-sap-cir cir-rate]
    - description description-string
    - no description
    - dest-mda [slot/mda | multipoint]
      - rate mda-rate
  - system
    - qos
      - access-ingress-aggregate-rate access-ingress-aggregate-rate [unshaped-sap-
cir cir-rate]
      - network-ingress-aggregate-rate network-ingress-aggregate-rate
```

9.3.1.2 Operational commands

```
config
- qos
  - copy fabric-profile src-prof dst-prof [overwrite]
```

9.3.1.3 Show commands

The **show>qos>fabric-profile** command applies only to the 7705 SAR-8 Shelf V2 and 7705 SAR-18; it does not apply to the fixed platforms. The **show>system>qos** command applies only to the fixed platforms, with the exception of the 7705 SAR-X. The 7705 SAR-X does not support fabric aggregate rate configuration; its shaper rate is set permanently to 16 Gb/s.

```
show
- qos
  - fabric-profile policy-id [association | detail]
show
- system
```

- qos

9.3.2 Command descriptions

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

9.3.2.1 Configuration commands

- [Generic commands](#)
- [QoS fabric profile commands](#)
- [Fixed platform fabric shaping commands](#)

9.3.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>fabric-profile

Description

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

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

Default

n/a

Parameters

description-string

a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

9.3.2.1.2 QoS fabric profile commands

fabric-profile

Syntax

```
fabric-profile policy-id [aggregate-mode | destination-mode] [create]  
no fabric-profile policy-id
```

Context

config>qos

Description

This command is used to configure a QoS fabric profile policy.

The default mode is **aggregate-mode**, which means that the access-ingress and network-ingress fabric shapers on all adapter cards are set to the same rate, either by default or via the **aggregate-rate** command. Selecting **destination-mode** allows each adapter card to have its fabric shapers set to a different rate for unicast traffic and a common rate for BMU traffic, either by default or via the **dest-mda rate** command.

The default fabric profile is preconfigured and non-modifiable. It cannot be deleted. All other fabric profiles must be explicitly created before use.

The create keyword must follow each new profile configuration.

Default

policy-id 1
aggregate-mode

Parameters

policy-id

the index identifier for a fabric profile policy

Values 1 to 256

aggregate-mode

assigns the aggregate fabric profile mode to the specified fabric profile

destination-mode

assigns the per-destination fabric profile mode to the specified fabric profile

create

keyword used to create a fabric profile policy

aggregate-rate

Syntax

aggregate-rate *aggregate-rate* [**unshaped-sap-cir** *cir-rate*]

Context

config>qos>fabric-profile

Description

This command sets the rate of the fabric shapers in aggregate mode, in kb/s. The **aggregate-rate** represents the maximum bandwidth that an adapter card can switch through its fabric interface. Each fabric shaper is set to the same aggregate rate. Using the **default** keyword sets the aggregate rate to 200 000 kb/s.

The **unshaped-sap-cir** command is optional. When used, it assigns an aggregate CIR to all unshaped 4-priority SAPs so that the unshaped SAP aggregate rate can compete with the shaped SAPs' aggregate CIR when the fabric scheduler selects traffic to forward to the switching fabric.

The default *cir-rate* value is 0 kb/s. If the *cir-rate* is **max**, then the maximum possible backplane speed for the adapter card, based on the associated fabric policy, is used.

The *aggregate-rate* can be higher or lower than the *cir-rate*. If the *cir-rate* is higher than the *aggregate-rate*, then the *cir-rate* is used by the network processor but the end result is that backpressure from the fabric shaper is applied indirectly to the per-SAP queues, as described below:

- per-SAP (aggregate) shapers (that is, unshaped SAP aggregate CIR or individual shaped SAP CIRs) prioritize traffic toward the fabric shaper based on the *cir-rate*
- the fabric shaper shapes the traffic at the configured rate
- when the fabric shaper becomes congested, it applies backpressure to individual SAP shapers; that is, it applies backpressure to the shaped SAPs and to the unshaped SAP aggregate, which in turn apply backpressure to individual unshaped queues

The same fabric-profile binding can be used for either network or access on a per-adapter card basis. However, the **unshaped-sap-cir** is not applicable when bound to network ingress.

Default

200000 ("**default**") for *aggregate-rate*

0 for *cir-rate*

Parameters

aggregate-rate

the rate of the fabric shapers in aggregate mode, in kb/s. Using the **default** keyword sets the aggregate rate to 200 000 kb/s.

Values

1 to 2500000 (2 500 000), or **default** for the 7705 SAR-8 Shelf V2 (all 6 slots)

1 to 10000000 (10 000 000), or **default** for the 7705 SAR-8 Shelf V2 (slots 1 and 2)

1 to 1000000 (1 000 000) or 1 to 2500000 ((2 500 000), or **default** for the 7705 SAR-18 (12 MDA slots)

1 to 10000000 (10 000 000), or **default** for the 7705 SAR-18 (4 XMDA slots)

cir-rate

the CIR rate for the unshaped SAP aggregate shapers, in kb/s. Using the **max** keyword sets the CIR rate to the maximum possible backplane speed for the adapter card, based on the associated fabric policy.

Values

1 to 2500000 (2 500 000), or **max** for the 7705 SAR-8 Shelf V2 (all 6 slots)

1 to 10000000 (10 000 000), or **max** for the 7705 SAR-8 Shelf V2 (slots 1 and 2)

1 to 1000000 (1 000 000) or 1 to 2500000 ((2 500 000), or **max** for the 7705 SAR-18 (12 MDA slots)

1 to 10000000 (10 000 000), or **max** for the 7705 SAR-18 (4 XMDA slots)

dest-md

Syntax

dest-md [*slot/mda* | **multipoint**]

Context

config>qos>fabric-profile

Description

This command enables the context for setting the rate for per-destination mode shapers on a specific adapter card or all adapter cards. Using the *slot/mda* parameter specifies a particular destination adapter card. Using the **multipoint** keyword specifies that all adapter cards are destination cards and they will have their rate configured to the same value. The value of the rate is configured using the **rate** command.

Parameters

slot/mda

the slot and mda identifier of the adapter card

Values slot: 1

mda: 1 to 6 on the 7705 SAR-8 Shelf V2, 1 to 12 on the 7705 SAR-18

multipoint

specifies that all adapter cards are destination adapter cards

rate

Syntax

rate *mda-rate*

Context

config>qos>fabric-profile>dest-md

Description

This command sets the rate of the fabric shapers in per-destination mode, in kb/s. When the **multipoint** keyword is used in the **dest-md** command, **rate** sets the bandwidth available to multipoint traffic through the fabric shapers; all adapter card shapers are set to the same value.

Using the **default** keyword sets the rate to 200 000 kb/s.

Default

200000 ("default")

Parameters

mda-rate

the rate of the fabric shapers in destination mode, in kb/s. Using the **default** keyword sets the *mda-rate* to 200 000 kb/s.

Values

- 1 to 2500000 (2 500 000), or **default** for the 7705 SAR-8 Shelf V2 (all 6 slots)
- 1 to 10000000 (10 000 000), or **default** for the 7705 SAR-8 Shelf V2 (slots 1 and 2)
- 1 to 1000000 (1 000 000) or 1 to 2500000 ((2 500 000), or **default** for the 7705 SAR-18 (12 MDA slots)
- 1 to 10000000 (10 000 000), or **default** for the 7705 SAR-18 (4 XMDA slots)

9.3.2.1.3 Fixed platform fabric shaping commands

access-ingress-aggregate-rate

Syntax

access-ingress-aggregate-rate *access-ingress-aggregate-rate* [**unshaped-sap-cir** *cir-rate*]

Context

config>system>qos

Description

This command sets the rate for fabric shapers on the 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M, or 7705 SAR-Wx for access ingress traffic. The rate represents the maximum bandwidth that the node can switch through its fabric interface. Only aggregate mode is supported.

This command is not supported on the 7705 SAR-X, where the shaper rate is set permanently to 16 Gb/s.

The **unshaped-sap-cir** command is optional. When used, it assigns a *cir-rate* to the unshaped SAP aggregate so that the unshaped SAP aggregate rate can compete with the shaped SAPs when the round-robin fabric scheduler selects traffic to forward to the switching fabric. When not used, the *cir-rate* is set to its default value of 0 kb/s. If the *cir-rate* is **max**, the maximum possible backplane speed is used.

The *access-ingress-aggregate-rate* can be higher or lower than the *cir-rate*. If the *cir-rate* is higher than the *access-ingress-aggregate-rate*, then the *cir-rate* is used by the network processor, but backpressure from the fabric shaper is applied indirectly to the per-SAP queues, as follows:

- per-SAP (aggregate) shapers (that is, unshaped SAP aggregate CIR or individual shaped SAP CIRs) prioritize traffic toward the fabric shaper based on the *cir-rate*
- the fabric shaper shapes the traffic at the configured rate
- when the fabric shaper becomes congested, it applies backpressure to individual SAP shapers; that is, it applies backpressure to shaped SAPs and to unshaped SAP aggregate, which in turn applies backpressure to individual unshaped queues

Default

500000 ("**default**") for *access-ingress-aggregate-rate*

0 for *cir-rate*

Parameters

access-ingress-aggregate-rate

the PIR rate of the fabric shapers for access ingress traffic, in kb/s. Using the **default** keyword sets the aggregate rate to 500 000 kb/s.

Values 1 to 10000000 (10 000 000), or **default**

cir-rate

the CIR rate of the unshaped SAP shaper for access ingress traffic, in kb/s. Using the **max** keyword sets the CIR rate to the maximum possible backplane speed.

Values 0 to 10000000 (10 000 000), or **max**



Note: The actual supported maximum rates depend on the platform:

- 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-M, and 7705 SAR-Wx: 5 Gb/s
- 7705 SAR-H: 4 Gb/s
- 7705 SAR-Hc: 2.5 Gb/s

network-ingress-aggregate-rate

Syntax

network-ingress-aggregate-rate *network-ingress-aggregate-rate*

Context

config>system>qos

Description

This command sets the rate for fabric shapers on the 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M, or 7705 SAR-Wx for network ingress traffic, in kb/s. The rate represents the maximum bandwidth that the node can switch through its fabric interface. Only aggregate mode is supported.

This command is not supported on the 7705 SAR-X, where the shaper rate is set permanently to 16 Gb/s.

Default

2000000 (**default**)

Parameters

network-ingress-aggregate-rate

the rate of the fabric shapers for network ingress traffic, in kb/s. Using the **default** keyword sets the aggregate rate to 2 000 000 kb/s (2 Gb/s).

Values 1 to 10000000 (10 000 000), or **default**



Note: The actual supported maximum rates depend on the platform:

- 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-M, and 7705 SAR-Wx: 5 Gb/s
- 7705 SAR-H: 4 Gb/s
- 7705 SAR-Hc: 2.5 Gb/s

9.3.2.2 Operational commands

copy

Syntax

copy **fabric-profile** *src-prof dst-prof* [**overwrite**]

Context

config>qos

Description

This command copies a source QoS fabric profile into a destination QoS fabric profile. If the destination profile was already defined, the keyword **overwrite** must be appended for the copy to complete.

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

Parameters

src-prof dst-prof

specifies the source profile ID that will be copied and the destination profile ID that the source fabric profile will be copied to

Values 1 to 256

overwrite

specifies that the existing destination profile is to be replaced. Everything in the existing destination profile will be overwritten with the contents of the source profile. If **overwrite** is not specified, an error will occur if the destination profile ID exists.

9.3.2.3 Show commands



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

fabric-profile

Syntax

fabric-profile [*policy-id*] [**association** | **detail**]

Context

show>qos

Description

This command displays QoS fabric profile information. If *policy-id* is not included in the command, the CLI shows a list of fabric policies.

Parameters

policy-id
specifies the QoS fabric profile
Values 1 to 256

association
displays all adapter cards to which the specified profile is assigned and in which direction; that is, network ingress to fabric or access ingress to fabric

detail
displays detailed policy information including policy associations

Output

The following output is an example of QoS fabric profile information, and [Table 62: QoS fabric profile field descriptions](#) describes the fields.

Output example

```
*A:ALU-1>show>qos# fabric-profile
=====
Fabric Profile
=====
Policy-Id  Mode      Description
-----
1          aggregate Default Fabric Profile QoS policy.
2          aggregate Fast Shaping fabric profile policy.
100        destination
=====
*A:ALU-1>show>qos# fabric-profile

*A:ALU-1>show>qos# fabric-profile 4 detail
=====
QoS Fabric Profile
=====
Fabric Profile Id (4)
-----
Policy-id   : 4
Mode        : destination
Description  : Fabric profile QoS policy 4
-----
Destination MDA      Rate (Kbps)
-----
1/1                 400000
1/2                 400000
1/3                 400000
1/4                 400000
1/5                 400000
1/6                 400000
multipoint          400000
-----
Associations
-----
```

```

MDA          : 1/1 (Network Ingress)
MDA          : 1/1 (Access Ingress)
MDA          : 1/2 (Network Ingress)
MDA          : 1/2 (Access Ingress)
MDA          : 1/3 (Network Ingress)
MDA          : 1/3 (Access Ingress)
MDA          : 1/4 (Network Ingress)
MDA          : 1/4 (Access Ingress)
MDA          : 1/5 (Network Ingress)
MDA          : 1/5 (Access Ingress)
MDA          : 1/6 (Network Ingress)
MDA          : 1/6 (Access Ingress)
=====
*A:ALU-1>show>qos#

*A:ALU-1>show>qos# fabric-profile 2 detail
=====
QoS Fabric Profile
=====
-----
Fabric Profile Id (2)
-----
Policy-id      : 2
Mode           : aggregate
Description     : Fast Shaping fabric profile policy

Aggregate-rate : 200 Kbps
Unshaped-sap-cir : 200 Kbps
-----
Associations
-----
MDA          : 1/1 (Network Ingress)
MDA          : 1/1 (Access Ingress)
MDA          : 1/2 (Network Ingress)
MDA          : 1/2 (Access Ingress)
MDA          : 1/3 (Network Ingress)
MDA          : 1/3 (Access Ingress)
MDA          : 1/4 (Network Ingress)
MDA          : 1/4 (Access Ingress)
MDA          : 1/5 (Network Ingress)
MDA          : 1/5 (Access Ingress)
MDA          : 1/6 (Network Ingress)
MDA          : 1/6 (Access Ingress)
=====

```

Table 62: QoS fabric profile field descriptions

Label	Description
Policy-id	The fabric profile QoS policy number
Mode	The fabric profile mode, either aggregate or destination
Description	The description associated with the fabric profile
Aggregate-rate	The rate of the fabric shapers in aggregate mode, in kb/s
Unshaped-sap-cir	The CIR rate of the aggregate SAP shaper for all unshaped 4-priority SAPs, in kb/s

Label	Description
Destination MDA/Rate (Kbps)	The slot/mda number of the destination adapter card, and the rate of the fabric shapers in per-destination mode, in kb/s
Associations	The adapter cards to which the specified fabric profile is assigned and in which direction, that is, network ingress to fabric or access ingress to fabric

qos

Syntax

qos

Context

show>system

Description

This command displays system QoS information for the 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M, and 7705 SAR-Wx.

This command is not supported on the 7705 SAR-X.

Output

The following output is an example of system QoS information, and [Table 63: System QoS field descriptions](#) describes the fields.

Output example

```
*A:ALU-1>show>system# qos
=====
System QoS Configuration
=====
Access Ingress Aggregate Rate      : 1234 Kbps
Access Ingress Unshaped Sap CIR    : 123 Kbps

Network Ingress Aggregate Rate     : 456 Kbps
=====
*A:ALU-1>show>system# qos
```

Table 63: System QoS field descriptions

Label	Description
Access Ingress Aggregate Rate	The rate of the access ingress fabric shapers in aggregate mode, in kb/s
Access Ingress Unshaped Sap CIR	The CIR rate of the access ingress unshaped SAP shapers in aggregate mode, in kb/s

Label	Description
Network Ingress Aggregate Rate	The rate of the network ingress fabric shapers in aggregate mode, in kb/s

10 QoS shapers and shaper QoS policies

This chapter provides information to configure shaper QoS policies using the command line interface.

Topics in this chapter include:

- [Overview](#)
- [Basic configuration](#)
- [Service management tasks](#)
- [Shaper QoS policy command reference](#)

10.1 Overview

Shapers are used to control access or network traffic flows through the 7705 SAR.

QoS shapers refer to per-SAP and per-VLAN dual-rate aggregate shapers that operate at the QoS tier 2 level.

Shaper QoS policies refer to policies that contain shaper groups, such as per-customer (MSS) dual-rate aggregate shapers that operate at the QoS tier 3 level.

Configuring a dual-rate aggregate shaper means setting PIR and CIR values. For per-customer (MSS) aggregate shapers, this is done through a shaper group.

Dual-rate aggregate shapers are applied to shaped and unshaped SAP traffic for access ingress and access egress flows, and to shaped and unshaped VLAN traffic for network egress flows.

On a hybrid port, where access and network traffic can share the same physical port, a shaper policy can be applied independently to access egress or network egress traffic.

For more information about QoS shapers and shaper QoS policies, see the following sections:

- [Per-SAP aggregate shapers \(H-QoS\) on Gen-2 hardware](#)
- [Per-VLAN network egress shapers](#)
- [Per-customer aggregate shapers \(multiservice site\) on Gen-2 hardware](#)
- [QoS for hybrid ports on Gen-2 hardware](#)
- [QoS for Gen-3 adapter cards and platforms](#)

10.2 Basic configuration

This section contains the following topics related to creating and applying shaper QoS policies:

- [Creating a shaper QoS policy and shaper groups](#)
- [Applying a shaper QoS policy and shaper groups](#)
- [Default shaper QoS policy values](#)
- [Configuring per-SAP aggregate shapers and an unshaped SAP aggregate shaper \(H-QoS\)](#)

- [Configuring per-VLAN shapers and an unshaped VLAN shaper](#)

A basic shaper QoS policy must conform to the following rules:

- Each shaper policy must have a unique policy identifier (name).
- Default values can be modified but parameters cannot be deleted.

10.2.1 Creating a shaper QoS policy and shaper groups

Configuring and applying shaper QoS policies is optional. If no shaper QoS policy is explicitly defined, the default shaper QoS policy ("default") is applied.

To create a new shaper policy, define the following:

- a shaper policy name – the system does not dynamically assign a name
- a description (optional) – a brief description of the policy
- a shaper group (optional) and its PIR and CIR values – a shaper group is a dual-rate aggregate shaper that is used mainly as a per-customer (MSS) aggregate shaper
- an unshaped SAP shaper group (optional) – the shaper group used by the group of all unshaped SAPs for the purpose of traffic arbitration with the shaped SAPs on a port

Use the following CLI syntax to configure a shaper policy and shaper groups within the shaper policy:

CLI syntax:

```
config>qos#
  shaper-policy policy-name [create]
    description description-string
    shaper-group shaper-group-name [create]
      description description-string
      rate pir-rate [cir cir-rate]
    unshaped-sap-shaper-group shaper-group-name
```

Example:

```
config>qos
config>qos$ shaper-policy "shaper_policy_2" create
config>qos>shaper-policy$ shaper-group "sg1"
config>qos>shaper-policy>shaper-group$ rate 1000 cir 200
config>qos>shaper-policy>shaper-group$ exit
config>qos>shaper-policy>shaper-group$ shaper-group "unshaped_sg1"
config>qos>shaper-policy>shaper-group$ rate 500 cir 100
config>qos>shaper-policy>shaper-group$ exit
config>qos>shaper-policy>unshaped-sap-shaper-group$ "unshaped_sg1"
config>qos>shaper-policy$ exit
config>qos$
```

The following output shows the configuration for "shaper_policy_2":

```
*A:7705custDoc:Sar18>config>qos# info detail
#-----
....
  shaper-policy "shaper_policy_2" create
    no description
    shaper-group default create
      description "Default Shaper Group."
      rate max cir 0
    exit
  shaper-group sg1 create
```

```

        description "Shaper Group_1"
        rate 1000 cir 200
    exit
    shaper-group unshaped_sg1 create
        no description
        rate 500 cir 100
    exit
    unshaped-sap-shaper-group unshaped_sg1
exit
.....
#-----
*A:7705custDoc:Sar18>config>qos#

```

10.2.2 Applying a shaper QoS policy and shaper groups

A shaper QoS policy must be assigned to an Ethernet MDA for access ingress per-customer aggregate shaping, or to a port for access egress per-customer aggregate shaping. After a shaper policy is assigned, a shaper group can be applied.

Shaper groups are created within the shaper policy and provide the per-customer (MSS) aggregate shapers. The unshaped SAP shaper group within the policy provides the shaper rate for all the unshaped SAPs (4-priority scheduled SAPs). For each shaped SAP, an ingress or egress shaper group can be specified.

For ingress, the shaper group assigned to a SAP or unshaped shaper group must be a shaper group from the shaper policy assigned to the Ethernet MDA.

For egress, the shaper group assigned to a SAP or unshaped shaper group must be a shaper group from the shaper policy assigned to the port.

All ingress shaped or unshaped SAPs configured with the same ingress shaper group on an Ethernet MDA have their aggregate traffic shaped at the shaper group rate. Similarly, all the egress shaped or unshaped SAPs configured with the same egress shaper group on a port have their aggregate traffic shaped at the shaper group rate.

On all 7705 SAR fixed platforms (with the exception of the 7705 SAR-X), when a shaper policy is assigned to an Ethernet MDA for access ingress aggregate shaping, it is automatically assigned to all the Ethernet MDAs in that chassis. The shaper group members contained in the shaper policy span all the Ethernet MDAs. SAPs on different Ethernet MDAs configured with the same ingress shaper group will share the shaper group rate.

For hybrid ports, one shaper policy can be applied to the network egress traffic as well as to the access egress traffic. For network egress, all the shaped interfaces (VLANs) and unshaped interfaces are bound to the default shaper group contained in the shaper policy assigned to the network egress port. The access egress traffic is bound to the shaper group assigned to each shaped SAP or to the default shaper group if none is assigned, and all the unshaped SAPs are assigned to the unshaped shaper group or to the default shaper group if none is assigned. Traffic is then scheduled between the network and access shaper groups by ensuring that the CIR rate for each shaper group is scheduled before any of the excess information rate (EIR) traffic from any of the shaper groups.

10.2.2.1 Applying a shaper policy

The following examples illustrate the CLI syntax to apply a shaper policy to an MDA and a hybrid port (access and network egress).

CLI syntax:

```
config>card>mda#
access
    ingress
        shaper-policy policy-name
```

CLI syntax:

```
config>port>ethernet#
egress-rate sub-rate
mode hybrid
access
    egress
        shaper-policy policy-name
        unshaped-sap-cir cir-rate
network
    egress
        shaper-policy policy-name
        unshaped-if-cir cir-rate
```

The following outputs show an ingress shaper policy applied to an MDA and different shaper policies applied to access egress and network egress traffic on a hybrid port:

```
*A:ALU>config>card>mda>access# info detail
-----
                ingress
                fabric-policy 1
                security-queue-policy 1
                shaper-policy "test_shaper_policy"
                exit
-----
*A:ALU>config>card>mda>access#
```

```
*A:ALU>config>port>ethernet# info
-----
                mode hybrid
                encap-type dot1q
                network
                egress
                shaper-policy 2
                unshaped-if-cir 250000
                exit
            exit
        access
        egress
        unshaped-sap-cir 200000
        shaper-policy 5
        exit
    exit
-----
```

10.2.2.2 Applying a shaper group

Shaper groups are applied to SAPs on Epipe, lpipe, and VPLS services, and to interface SAPs on IES and VPRN services. The shaper group must exist within the shaper policy assigned to an MDA.

When an **unshaped-sap-shaper-group** is configured within a shaper policy, it is automatically applied to the unshaped SAPs on the MDA. Operators do not need to specifically apply the unshaped SAP shaper group.

Use the following CLI syntax to apply a shaper group to a VPLS SAP. The syntax is similar for Epipe and lpipe SAPs. The syntax is also similar for IES and VPRN services, except that the SAP is configured on a service interface.

CLI syntax:

```
config>service>vpls>sap#
  egress
    [no] shaper-group name
  ingress
    [no] shaper-group name
```

The following output shows shaper group "sg1" applied to access ingress traffic on a VPLS SAP:

```
*A:ALU>config>service>vpls>sap>ingress# info detail
-----
                qos 1
                no match-qinq-dot1p
                scheduler-mode 16-priority
                no agg-rate-limit
                shaper-group sg1
                -----
*A:ALU>config>service>vpls>sap>ingress#
```

10.2.2.3 Viewing shaper policy information

Use the **show qos shaper-policy** command to view information about a shaper policy, its shaper groups (including the unshaped SAP shaper group), the PIR and CIR values, and where the policy is used.

```
*A:ALU>show qos shaper-policy "test_shaper_policy" detail
=====
QoS Shaper Policy
=====
-----
Shaper Policy (test_shaper_policy)
-----
Policy                : test_shaper_policy
Description            : (Not Specified)
Unshaped Sap Shaper Group : unshaperd_sg1
-----
-----
Shaper Group Name      PIR (Kbps)    CIR (Kbps)
-----
default                max            0
test_sg1               999000       555000
unshaperd_sg1         888000       444000
-----
-----
Policy Associations
```

```

-----
Object Type      Object Id      Direction
-----
MDA              1/1           access ingress
MDA              1/2           access ingress
MDA              1/5           access ingress
Port             1/2/1         access egress
Port             1/5/8         network egress
-----
=====
*A:ALU>show qos shaper-policy "test_shaper_policy" detail#

```

10.2.3 Default shaper QoS policy values

The default shaper policies are identified by the policy name "default". The default policies cannot be deleted. The following table displays the default shaper policy parameters.

Table 64: Shaper policy defaults

Field	Default
description	"Default Shaper QoS policy."
shaper-group	"default"
description	"Default Shaper Group."
pir-rate	max
cir-rate	0

The following output displays the default configuration:

```

A*A:ALU-1>config>qos# info detail
#-----
echo "QoS Policy Configuration"
#-----
...
    shaper-policy "default" create
        description "Default Shaper QoS policy."
    shaper-group default create
        description "Default Shaper Group."
        rate max cir 0
    exit
exit
...
#-----

```

10.2.4 Configuring per-SAP aggregate shapers and an unshaped SAP aggregate shaper (H-QoS)

Configuring tier 2 per-SAP aggregate shapers applies to Gen-2 SAPs configured with 16-priority scheduling mode and Gen-3 SAPs. Access ingress and access egress shapers are configured under the same CLI context.

Configuring a single tier 2 aggregate shaper for all unshaped SAPs applies to Gen-2 SAPs configured with 4-priority scheduling mode. Access ingress and access egress shapers are configured under different CLI contexts:

- [Creating 16-priority shaped SAPs and configuring per-SAP aggregate shapers](#)
- [Configuring an unshaped aggregate CIR for all 4-priority unshaped SAPs \(access ingress\)](#)
- [Configuring an unshaped aggregate CIR for all 4-priority unshaped SAPs \(access egress\)](#)

10.2.4.1 Creating 16-priority shaped SAPs and configuring per-SAP aggregate shapers

Create shaped Gen-2 SAPs by configuring the scheduler mode to 16-priority. On Gen-3 hardware, the scheduler mode is always 4-priority and SAPs are shaped.

Configure a shaped SAPs' aggregate rates by setting the CIR and PIR, as required. The SAP must be shut down before the scheduler mode and the *agg-rate* and *cir-rate* can be changed.

Use the first CLI syntax (below) to create shaped SAPs for access egress and access ingress for VLL and VPLS services (including routed VPLS), and the second syntax for IES and VPRN services. Examples are provided for Epipe and IES services on Gen-2 and Gen-3 hardware:

CLI syntax:

```
config>service>epipe service-id customer customer-id create
  sap sap-id create
    egress
      scheduler-mode {4-priority | 16-priority}
      agg-rate-limit agg-rate [cir cir-rate]
    ingress
      scheduler-mode {4-priority | 16-priority}
      agg-rate-limit agg-rate [cir cir-rate]
```

CLI syntax:

```
config>service>ies service-id customer customer-id create
  interface ip-interface-name create
  sap sap-id create
    egress
      scheduler-mode {4-priority | 16-priority}
      agg-rate-limit agg-rate [cir cir-rate]
    ingress
      scheduler-mode {4-priority | 16-priority}
      agg-rate-limit agg-rate [cir cir-rate]
```

In the examples below, MDA 1/12 is a Gen-2 adapter card and MDA 1/3 is a Gen-3 adapter card.

Example:

```
config>service>epipe# sap 1/12/8 create
config>service>epipe>sap# shutdown
config>service>epipe>sap# egress
  ...egress# scheduler-mode 16-priority
  ...egress# agg-rate-limit 250000 cir 150000
config>service>epipe# sap 1/12/7 create
config>service>epipe>sap# ingress
  ...ingress# scheduler-mode 16-priority
  ...ingress# agg-rate-limit 250000 cir 150000
config>service>epipe# sap 1/3/2 create
config>service>epipe>sap# ingress
```

```
...ingress# agg-rate-limit 250000 cir 150000
```

Example:

```
config>service>ies# interface "ies_6000_interface"
config>service>ies>if# sap 1/12/5 create
config>service>ies>if>sap# shutdown
config>service>ies>if>sap# egress
...egress# scheduler-mode 16-priority
...egress# agg-rate-limit 250000 cir 150000
config>service>ies# interface "ies_6001_interface"
config>service>ies>if# sap 1/12/6 create
config>service>ies>if>sap# # ingress
...ingress# scheduler-mode 16-priority
...ingress# agg-rate-limit 250000 cir 150000
config>service>ies>if# sap 1/3/3 create
config>service>ies>if>sap# # ingress
...ingress# agg-rate-limit 250000 cir 150000
```

The following outputs display the shaped SAP configuration for Epipe and IES services:

```
*A:7705custDoc:Sar18>config>service>epipe# info
-----
....
    sap 1/12/7 create
        shutdown
        ingress
            scheduler-mode 16-priority
            agg-rate-limit 250000 cir 150000
        exit
    exit
    sap 1/12/8 create
        shutdown
        egress
            scheduler-mode 16-priority
            agg-rate-limit 250000 cir 150000
        exit
    exit
    sap 1/3/2 create
        shutdown
        egress
            agg-rate-limit 250000 cir 150000
        exit
    exit
-----
*A:7705custDoc:Sar18>config>service>epipe#
```

```
*A:7705custDoc:Sar18>config>service>ies# info
-----
....
    interface "ies_6000_interface" create
        sap 1/12/5 create
            shutdown
            egress
                scheduler-mode 16-priority
                agg-rate-limit 250000 cir 150000
            exit
        exit
    exit
    interface "ies_6001_interface" create
        sap 1/12/6 create
            shutdown
```



```

        ingress
            scheduler-mode 16-priority
            agg-rate-limit 250000 cir 150000
        exit
    exit
exit
interface "ies_6002_interface" create
sap 1/3/3 create
shutdown
egress
    agg-rate-limit 250000 cir 150000
exit
exit
-----

```

10.2.4.2 Configuring an unshaped aggregate CIR for all 4-priority unshaped SAPs (access ingress)

Use the first CLI syntax to set the **unshaped-sap-cir** for all unshaped 4-priority SAPs on a 7705 SAR-8 Shelf V2 or 7705 SAR-18. The fabric profile must be in aggregate mode before an unshaped SAP aggregate CIR can be assigned. Use the second CLI syntax for a 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, or 7705 SAR-Wx; it does not apply to the 7705 SAR-X.

CLI syntax:

```

config>qos
    fabric-profile policy-id aggregate-mode create
        aggregate-rate aggregate-rate [unshaped-sap-cir cir-rate]

```

CLI syntax:

```

config>system>qos
    access-ingress-aggregate-rate access-ingress-aggregate-rate [unshaped-sap-cir cir-rate]

```

Example:

```

config>qos#
config>qos# fabric-profile 2 aggregate-mode create
config>qos>fabric-profile# aggregate-rate 250000 unshaped-sap-cir 150000

```

Example:

```

config>system>qos#
config>system>qos# access-ingress-aggregate-rate 250000 unshaped-sap-cir 150000

```

The following output displays the fabric profile configuration for an H-QoS configuration:

```

*A:ALU-1:Sar18>config>qos# info detail
#-----
echo "QoS Policy Configuration"
#-----
....
    fabric-profile 2 aggregate-mode create
        no description
        aggregate-rate 250000 unshaped-sap-cir 150000
    exit
#-----

```

```
*A:ALU-1>config>qos#
```

For the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, and 7705 SAR-Wx:

```
*A:7705custDoc:sarM>config>system>qos# info detail
-----
      access-ingress-aggregate-rate 250000 unshaped-sap-cir 150000
      network-ingress-aggregate-rate 2
-----
*A:7705custDoc:sarM>config>system>qos#
```

10.2.4.3 Configuring an unshaped aggregate CIR for all 4-priority unshaped SAPs (access egress)

To provide arbitration between the bulk (aggregate) of unshaped SAPs and the shaped SAPs, assign a rate to the unshaped SAPs. See the 7705 SAR Interface Configuration Guide for command descriptions.

Use the following CLI syntax to configure a per-port CIR rate limit for the aggregate of all 4-priority unshaped SAPs on the egress port:

CLI syntax:

```
config>port port-id>ethernet
      access
      egress
      [no] unshaped-sap-cir cir-rate
```

Example:

```
config# port 1/12/7
config>port# ethernet
config>port>ethernet# access
config>port>ethernet>access# egress
config>port>ethernet>access>egress# unshaped-sap-cir 5000000
```

The following output displays the port configuration for the unshaped SAPs in an H-QoS configuration:

```
*A:7705custDoc:Sar18>config>port# info
-----
.....
      ethernet
      access
      egress
      unshaped-sap-cir 500000
      exit
      exit
      exit
-----
*A:7705custDoc:Sar18>config>port#
```

10.2.5 Configuring per-VLAN shapers and an unshaped VLAN shaper

This section contains information about the following topics:

- [Configuring per-VLAN network egress shapers](#)
- [Configuring a CIR for network egress unshaped VLANs](#)

10.2.5.1 Configuring per-VLAN network egress shapers

Per-VLAN network egress shapers can be configured for network interfaces.

The **queue-policy** command is used to enable and disable the network egress per-VLAN shapers on a per-interface basis. If the **no queue-policy** command is used, the VLAN (that is, the interface) defaults to unshaped mode. The **agg-rate-limit** command cannot be accessed unless a network queue policy is assigned to the interface.

Use the following CLI syntax to configure a per-VLAN network egress shaper on a network interface. See the 7705 SAR Router Configuration Guide for command descriptions.

CLI syntax:

```
config>router>interface#
    egress
        queue-policy name
        agg-rate-limit agg-rate [cir cir-rate]
```

10.2.5.2 Configuring a CIR for network egress unshaped VLANs

To provide arbitration between the bulk (aggregate) of unshaped VLANs and the shaped VLANs, assign a rate to the unshaped VLANs.

Use the following CLI syntax to configure a CIR for the bulk of network egress unshaped VLANs. See the 7705 SAR Interface Configuration Guide for command descriptions.

CLI syntax:

```
config>port>ethernet#
    network
        egress
            unshaped-if-cir cir-rate
```

10.3 Service management tasks

This section describes the following service management tasks:

- [Removing and deleting QoS policies](#)
- [Copying and overwriting QoS policies](#)
- [Editing QoS policies](#)

10.3.1 Removing and deleting QoS policies

A QoS policy cannot be deleted until its associations with all the ports are removed.

Because one shaper policy can be assigned to multiple ports, you must remove all the associations to the ports before deleting the shaper policy.

Use the following CLI syntax to remove a shaper policy from a port and then delete the shaper policy from the QoS configuration:

CLI syntax:

```
config>port>ethernet>access>egress# no shaper-policy
config>qos# no shaper-policy policy-name
```

Example:

```
config>port>ethernet>access>egress# no shaper-policy
config>qos# no shaper-policy shaper_policy_2
```

10.3.2 Copying and overwriting QoS policies

You can copy an existing shaper policy to a new shaper policy or overwrite an existing shaper policy. If the destination policy ID exists, the **overwrite** option must be specified or an error occurs.

Use the following syntax to overwrite an existing shaper QoS policy.

CLI syntax:

```
config>qos# copy shaper-policy src-name dst-name overwrite
```

Example:

```
config>qos# copy shaper-policy ShaperPolicy1 ShaperPolicy2 overwrite
config>qos# exit
```

The following output displays the copied policies:

```
*A:ALU-2>config>qos# info detail
#-----
echo "QoS Policy Configuration"
#-----
...
    shaper-policy "default" create
        description "Default shaper policy"
        shaper-group "default" create
            rate max cir 0
        exit
    exit
    shaper-policy "ShaperPolicy2" create
        description "Test2"
        shaper-group "default" create
            rate 2000 cir 250
        exit
    exit
    shaper-policy "ShaperPolicy1" create
        description "Test1"
        shaper-group "default" create
            rate 1000 cir 150
        exit
    exit
    shaper-policy "ShaperPolicy2" create
        description "Test1"
        shaper-group "default" create
            rate 1000 cir 150
        exit
    exit
...
```

```
#-----
```

10.3.3 Editing QoS policies

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

10.4 Shaper QoS policy command reference

10.4.1 Command hierarchies

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

10.4.1.1 Configuration commands

```
config
- qos
  - shaper-policy policy-name [create]
  - no shaper-policy policy-name
    - description description-string
    - no description
  - shaper-group shaper-group-name [create]
  - no shaper-group shaper-group-name
    - rate pir-rate [cir cir-rate]
    - no rate
  - [no] unshaped-sap-shaper-group shaper-group-name
```

10.4.1.2 Operational commands

```
config
- qos
  - copy shaper-policy src-name dst-name [overwrite]
```

10.4.1.3 Show commands

```
show
- qos
  - shaper-policy [shaper-policy-name] [detail]
```

10.4.2 Command descriptions

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

10.4.2.1 Configuration commands

- [Generic commands](#)
- [Shaper QoS policy commands](#)

10.4.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>qos>shaper-policy

Description

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

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

Default

n/a

Parameters

description-string

a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

10.4.2.1.2 Shaper QoS policy commands

shaper-policy

Syntax

```
shaper-policy policy-name [create]
no shaper-policy policy-name
```

Context

```
config>qos
```

Description

This command enables the context to configure a shaper QoS policy.

For hybrid ports, the shaper policy is independently assigned to access or network egress traffic. When the Ethernet port mode is changed to hybrid mode, the default policy is assigned to access and network traffic. To change an access or network policy, use the commands **config>port>ethernet>access>egress>shaper-policy** and **config>port>ethernet>network>egress>shaper-policy**.

For access ingress per-customer aggregate shaping, the shaper policy is assigned to an Ethernet MDA, and SAPs on that Ethernet MDA must be bound to a shaper group within the shaper policy bound to that Ethernet MDA. To assign a shaper policy to an adapter card, use the command **config>card>mda>access>ingress>shaper-policy**.

For access egress per-customer aggregate shaping, the shaper policy is assigned to a port, and SAPs on that port must be bound to a shaper group within the shaper policy bound to that port. To assign a shaper policy on egress, use the command **config>port>ethernet> access>egress>shaper-policy**.

The default shaper policy cannot be deleted. The following table displays the default shaper policy parameters.

Table 65: Shaper policy defaults

Field	Default
description	"Default Shaper QoS policy."
shaper-group	"default"
description	"Default Shaper Group."
pir-rate	max
cir-rate	0

The **no** form of this command removes the configured **shaper-policy**.

Default

shaper-policy "default"

Parameters

policy-name

the name of the shaper policy. To access the default shaper policy, enter "default".

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters.

If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

create

keyword used to create a shaper policy

shaper-group

Syntax

shaper-group *shaper-group-name* [**create**]

no shaper-group *shaper-group-name*

Context

config>qos>shaper-policy

Description

This command creates and configures a shaper group. A shaper group is a dual-rate aggregate shaper used to arbitrate shaped and unshaped traffic for SAPs (access) and for VLANs (network).

The default shaper group cannot be deleted.

The **no** form of this command removes the configured **shaper-group**.

Default

shaper-group "default"

Parameters

shaper-group-name

the name of the shaper group. To access the default shaper group, enter "default".

create

keyword used to create a shaper group

rate

Syntax

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

no rate

Context

config>qos>shaper-policy>shaper-group

Description

This command sets the PIR and CIR for the shaper group. When the PIR or CIR is set to **max**, the corresponding aggregate shaper rate is capped at the **egress-rate** configured on the port (**config>port>ethernet>egress-rate**).

The **no** form of this command restores the *pir-rate* and *cir-rate* to the default values.

Default

pir-rate: max cir-rate: 0

Parameters

cir-rate

the CIR for the shaper group

Values 0 to 1000000000, or max (kbps)

pir-rate

the PIR for the shaper group

Values 1 to 1000000000, or max (kbps)

unshaped-sap-shaper-group

Syntax

[**no**] **unshaped-sap-shaper-group** *shaper-group-name*

Context

config>qos>shaper-policy

Description

This command assigns a shaper group to the unshaped SAPs assigned to the shaper policy. An unshaped shaper group is a dual-rate aggregate shaper used to arbitrate unshaped traffic for SAPs (access).

There can be only one unshaped SAP shaper group per shaper policy. The unshaped SAP shaper group must already exist within the shaper policy before it can be chosen for the **unshaped-sap-shaper-group**.

The **no** form of this command removes the configured **unshaped-sap-shaper-group**.

Default

shaper-group "default"

Parameters

shaper-group-name

the name of the unshaped SAP shaper group. To access the default shaper group, enter "default".

10.4.2.2 Operational commands

copy

Syntax

copy shaper-policy *src-name dst-name* [**overwrite**]

Context

config>qos

Description

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

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

Parameters

shaper-policy *src-name dst-name*

indicates that the source policy ID and the destination policy ID are shaper 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 that the existing destination policy is to be replaced. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, an error will occur.

10.4.2.3 Show commands



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

shaper-policy

Syntax

shaper-policy [*shaper-policy-name*] [**detail**]

Context

show>qos

Description

This command displays shaper policy information.

Parameters

- shaper-policy-name*
the name of the shaper policy
- detail**
displays detailed information about the shaper policy

Output

The following output is an example of shaper policy information, and [Table 66: Shaper policy field descriptions](#) describes the fields.

Output example

```
*A:7705custDoc:Sar18>show>qos# shaper-policy
=====
Shaper Policies
=====
Policy-Id          Description
-----
2
5
default            Default Shaper QoS policy.
shaper_policy_1
shaper_policy_2
=====
*A:7705custDoc:Sar18>show>qos#

*A:7705:Dut-C# show qos shaper-policy 5
=====
QoS Shaper Policy
=====
-----
Shaper Policy (5)
-----
Policy              : 5
Description         : Description for Shaper Policy id # 5
-----
Shaper Group Name   PIR      CIR
-----
default             max      0
```

```

=====
*A:7705:Dut-C# show qos shaper-policy 5 detail
=====
QoS Shaper Policy
=====
-----
Shaper Policy (5)
-----
Policy                : 5
Description           : Description for Shaper Policy id # 5
-----
Shaper Group Name      PIR      CIR
-----
default                max      0
-----
Policy Associations
-----
Object Type   Object Id   Direction
-----
Port          1/10/7      access egress
Port          1/10/8      network egress
=====
*A:7705custDoc:Sar18>show>qos#

```

Table 66: Shaper policy field descriptions

Label	Description
Policy	The ID that uniquely identifies the policy
Description	A text string that helps identify the policy's context in the configuration file
Shaper Group Name	The name of the shaper group
PIR	The peak information rate for the shaper
CIR	The committed information rate for the shaper
Policy Associations	
Object Type	The type of object using the specified shaper policy
Object Id	The identifier of the object using the specified shaper policy
Direction	The direction of traffic to which the shaper policy applies

11 Security QoS and security QoS policies

This chapter provides information about security QoS used to control firewall traffic that is extracted to the CSM for examination. It also provides information about configuring security queue QoS policies using the command line interface.

Topics in this chapter include:

- [Overview](#)
- [QoS for firewall-extracted packets to the CSM](#)
- [Multi-chassis firewall QoS](#)
- [Security queue QoS policies](#)
- [Basic configuration](#)
- [Service management tasks](#)
- [Security queue QoS policy command reference](#)

11.1 Overview

When a security zone, security profile, and policies are configured for security sessions on the 7705 SAR, data packets entering and leaving the zone are extracted, if required, from the datapath to the CSM for examination. QoS is applied on these packets to control the amount of traffic being extracted to the CSM. For information about requirements for packet extraction to the CSM, see the "Security Session Creation" in the 7705 SAR Router Configuration Guide.

11.2 QoS for firewall-extracted packets to the CSM

When security parameters are configured, data packets entering or leaving a zone are extracted from the datapath to the CSM for examination. Application Level Gateway (ALG) TFTP/FTP or strict TCP data packets that are extracted are placed into access or network security data queues. These access and network security queues are able to control the rate of traffic scheduled through these queues by using security queue QoS policies (see [Security queue QoS policies](#) for information).

Non-ALG and non-strict TCP datapath traffic that is extracted from the datapath for CSM security examination is extracted into a security control queue that has one queue per security zone.

To limit the aggregate datapath traffic being extracted to the CSM via the access/network security queues and all the security control queues (one per zone), a **security-aggregate-rate** shaper can be configured, which defaults to a rate of 50 Mb/s. For information about configuring the **security-aggregate-rate** shaper, see the 7705 SAR Interface Configuration Guide, "Adapter Card Commands".

Firewall traffic that is permitted through the firewall will be forwarded across the data path using datapath traffic management.

11.3 Multi-chassis firewall QoS

In a multi-chassis configuration, the slave router has the same security configuration as the master. When the slave router receives datapath packets that are entering or leaving a security zone, the data packets are extracted into the same access or network data queues and security control queues that exist on the master. However, the data packets that are extracted must be processed by the master firewall security engine. The slave sends these extracted data packets to the master over the multi-chassis link (MCL).

The access queues, network data queues, and security control queues used on the slave have QoS configurations that control the traffic rate from the slave to the master. These QoS configurations on the slave, specifically security queue QoS policies and the aggregate shaping rate, should be configured identically on the master. For information, see [Security queue QoS policies](#) and also refer to the 7705 SAR Interface Configuration Guide, "Adapter card commands" for information about configuring the **security-aggregate-rate** command.

The extracted data packets that the master receives from the slave are stored in a multi-chassis firewall queue for extraction to the CSM on the master. To limit the rate of datapath traffic being extracted and sent to the master CSM, this extraction queue is rate-limited to 80 Mb/s. In addition, this extraction queue, along with the security control queues and the access/network security queues, are rate-limited by the **security-aggregate-rate** command. These QoS settings and configurations make it possible to control the datapath traffic being extracted on the master and slave for firewall security processing.

11.4 Security queue QoS policies

For ALG TFTP/FTP or strict TCP traffic that egresses one security zone and ingresses a different security zone, every packet must be forwarded to the CSM for processing. To control this traffic to the CSM, the packets are extracted from the data path and queued into either network security data queues or access security data queues. These queues each contain two further queues: expedited (EXP) queues and best-effort (BE) queues. On the 7705 SAR-8 Shelf V2 and 7705 SAR-18, expedited and best-effort queues are created per adapter card.

For further details about zone configuration and firewall session creation, see the 7705 SAR Router Configuration Guide, "Configuring Security Parameters".

11.4.1 Packet queuing with DSCP

By default, packets are assigned to the EXP and BE queues as follows:

- For the base router context, packets are assigned to the EXP and BE queues based on the DSCP marking in the packet IP header.
- For the VPRN or IPsec context, packets are assigned to the EXP and BE queues based on the EXP or DSCP marking of the outer tunnel. The EXP marking is used for Layer 3 MPLS VPRNs, and the DSCP marking is used for IPsec or Layer 3 GRE VPRNs.

However, it is possible to queue packets based on the inner (customer) IP header DSCP marking by using the command **config>qos>network>ingress>ler-use-dscp**. This is useful where customers have policed bandwidth at the PE and want to differentiate their own network packets on the access PEs. By enabling the **ler-use-dscp** command, the following occurs for encrypted VPRN, IPsec, and NGE packets:

- packets will be queued in the encryption queues based on the outer tunnel MPLS EXP or IPSec/GRE DSCP marking
- after decryption, for either firewall datapath queues or the regular datapath queues, the packets will be queued based on the inner (customer) IP header DSCP marking

For more information, see [ler-use-dscp](#) in the Network QoS Policy Command Reference chapter.

11.5 Basic configuration

This section contains the following topics related to creating security queue policies:

- [Creating a security data queue QoS policy](#)
- [Default security queue policy parameter values](#)

A basic security queue policy must conform to the following rules:

- Each security queue policy must have a unique policy ID.
- Default values can be modified but parameters cannot be deleted.



Note: Queue 1 is always best effort and queue 2 is always expedited.

11.5.1 Creating a security data queue QoS policy

Configuring a security data queue QoS policy is optional. If no security queue QoS policy is explicitly defined, the default security queue QoS parameters are applied.

To create a new security queue policy, define the following:

- a security queue policy identifier – the system does not dynamically assign an identifier
- a description – a brief description of the policy

Use the following CLI syntax to configure a security queue QoS policy:

CLI syntax:

```
config>qos#
  security-queue policy-id
    description description-string
    queue queue-id
      cbs size
      high-prio-only percent
      mbs size
      rate pir [cir]
```

Example:

```
*A:ALU-1#
config>qos>security-queue "SecurityQueue 2" create
config>qos>security-queue$ description "Test1"
config>qos>security-queue$ queue 1
config>qos>security-queue>queue$ cbs 112
config>qos>security-queue>queue$ high-prio-only 25
config>qos>security-queue>queue$ mbs 300 kilobytes
config>qos>security-queue>queue$ rate pir max cir max
config>qos>security-queue>queue$ exit
config>qos>security-queue$ queue 2
```



```

config>qos>security-queue>queue$c bs 40
config>qos>security-queue>queue$ mbs 5000
config>qos>security-queue>queue$ rate pir 400000 cir 35000
config>qos>security-queue>queue$ exit
config>qos>security-queue$ exit
*A:ALU-1#

```

The following output shows the configuration for SecurityQueue 2:

```

*A:ALU-1>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
    "SecurityQueue 2" create
        description "Test1"
        queue 1 best-effort
            rate max cir max
            mbs 300 kilobytes
            cbs 112
            high-prio-only 25
        exit
        queue 2 expedite
            rate 400000 cir 35000
            mbs 5000 kilobytes
            cbs 40
        exit
    exit
#-----

```

11.5.2 Default security queue policy parameter values

The following table displays the default security queue policy parameter values.

Table 67: Security queue parameter defaults

Parameter	Default values–Best Effort	Default values–Expedited
CBS	10 kB	40 kB
High-prio-only	10	n/a
MBS	5000 kB	5000 kB
PIR	400000 kB	400000 kB
CIR	1500 kB	35000 kB

11.6 Service management tasks

This section describes the following service management tasks:

- [Deleting QoS policies](#)
- [Copying and overwriting QoS policies](#)

- [Editing QoS policies](#)

11.6.1 Deleting QoS policies

Use the following CLI syntax to delete a security queue QoS policy:

CLI syntax:

```
config>qos# no security-queue policy-id
```

Example:

```
config>qos# no security-queue SecurityQueue 2
```

11.6.2 Copying and overwriting QoS policies

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

Use the following syntax to overwrite an existing security queue QoS policy.

CLI syntax:

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

Example:

```
*A:ALU-1>config>qos# copy security-queue SecurityQueue1 SecurityQueue2
overwrite
config>qos# exit
*A:ALU-2#
```

11.6.3 Editing QoS policies

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

11.7 Security queue QoS policy command reference

11.7.1 Command hierarchies

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

11.7.1.1 Configuration commands

```
config
- qos
  - security-queue policy-id [create]
  - no security-queue policy-id
    - description description-string
    - no description
  - [no] queue queue-id
    - cbs {size-in-kbytes | default}
    - no cbs
    - high-prio-only {percent | default}
    - no high-prio-only
    - mbs {size {bytes | kbytes} | default}
    - no mbs
    - rate pir [cir cir]
    - no rate
```

11.7.1.2 Operational commands

```
config
- qos
  - copy security-queue src-pol dst-pol [overwrite]
```

11.7.1.3 Show commands

```
show
- qos
  - security-queue [policy-id] [association | detail]
```

11.7.2 Command descriptions

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

11.7.2.1 Configuration commands

- [Security queue QoS policy commands](#)

11.7.2.1.1 Security queue QoS policy commands

security-queue

Syntax

security-queue *policy-id* [**create**]
no security-queue *policy-id*

Context

config>qos

Description

This command configures a security queue policy for traffic being extracted from the datapath to the CSM for firewall processing. When a security queue policy is created, two queues are created automatically for the extracted traffic: queue 1 for best-effort traffic and queue 2 for expedited traffic. The queue number and type for these two queues is not configurable.

The **no** form of this command removes the security queue policy.

Default

n/a

Parameters

- policy-id*

the number of the policy being referenced. Policy 1 is reserved for the default security queue policy; it cannot be modified.

Values 1 to 65535
- create**

keyword used to create a security queue policy

description

Syntax

description *description-string*

no description

Context

config>qos>security-queue

Description

This command configures a description for the security queue policy being referenced.

The **no** form of this command removes the description.

Default

n/a

Parameters

description-string

a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

queue

Syntax

[no] queue *queue-id*

Context

config>qos>security-queue

Description

This command enables the context to configure parameters related to the queue type for the traffic extracted from the datapath to the CSM. When the security queue policy is created, a set of queues is automatically created: queue 1 for best-effort traffic and queue 2 for expedited traffic. When the best-effort and expedited queues are created, default values are assigned to their information rate parameters.

The **no** form of this command removes the *queue-id* from the security queue policy.

Default

n/a

Parameters

<i>queue-id</i>	specifies the ID for the queue type being referenced
Values	1 for best effort queue
Values	2 for expedited queue

cbs

Syntax

cbs {*size-in-kbytes* | **default**}

no cbs

Context

config>qos>security-queue>queue

Description

This command overrides the default committed buffer space (CBS) reserved for the specified queue. The value is configured in kilobytes.

The **no** form of this command returns the CBS to the default value for the queue type.

Parameters

<i>size-in-kbytes</i>	specifies the committed buffer space for the queue
Values	1 to 131072 default
Default	10 kB for best effort 40 kB for expedite

high-prio-only

Syntax

high-prio-only {*percent* | **default**}

no high-prio-only

Context

config>qos>security-queue>queue

Description

This command configures the percentage of the queue used exclusively by high-priority packets. The specified value overrides the default value for the queue type.

The **no** form of this command restores the default high-priority reserved size for the queue type.

Parameters

percent
the percentage reserved for high priority traffic on the queue

Values	1 to 100 default
Default	10 for best effort 10 for expedite

mbs

Syntax

mbs {size {bytes | kilobytes} | default}
no mbs

Context

config>qos>security-queue>queue

Description

This command sets the maximum burst size (MBS) value for buffers of a specified queue. The value is configured either in bytes or in kilobytes and overrides the default MBS value.

The **no** form of this command returns the MBS to the default value for the queue type.

Parameters

size
specifies the maximum burst size for the queue, either in bytes or kilobytes

Values	0 to 131072000 default
Default	5000 kB for best effort 5000 kB for expedite

bytes
configures the maximum burst size for the queue in bytes

kilobytes
configures the maximum burst size for the queue in kilobytes

rate

Syntax

`rate pir [cir cir]`
`no rate`

Context

`config>qos>security-queue>queue`

Description

This command sets the peak information rate (PIR) value and optional committed information rate (CIR) for a specified queue. The values are configured in kilobytes and override the default PIR and CIR values.

The **no** form of this command returns the PIR and CIR to their default values for the queue type, assigned when the security queue policy for firewall traffic was created.

Parameters

<i>pir</i>	specifies the peak information rate for the queue, in kilobytes per second
Values	1 to 100000000 max
Default	400000 for best effort 400000 for expedite
<i>cir</i>	specifies the committed information rate for the queue, in kilobytes per second
Values	0 to 100000000 max
Default	15000 for best effort 35000 for expedite

11.7.2.2 Operational commands

copy

Syntax

`copy security-queue src-pol dst-pol [overwrite]`

Context

`config>qos`

Description

This command copies existing policy entries for a security queue QoS policy to another security queue policy. This command is a configuration-level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Default

n/a

Parameters

src-pol

the source policy ID that the copy command will attempt to copy from

dst-pol

the destination policy ID to which the command will copy the policy

overwrite

specifies that the existing destination policy is to be replaced. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified for an existing policy ID, an error will occur.

11.7.2.3 Show commands



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

security-queue

Syntax

security-queue [*policy-id*] [**association** | **detail**]

Context

show>qos

Description

This command displays security queue information.

Parameters

policy-id

specifies the ID of the security queue policy

Values 1 to 65535

association

displays information about the security queue policy associations

detail

displays detailed information about the security queue policy

Output

The following output is an example of security policy information, and [Table 68: Security policy field descriptions](#) describes the fields.

Output example

```
*A:7705custDoc:Sar18>show>qos# security-queue detail
=====
QoS Security Queue Policy
=====
Security Queue Policy Id (1)
-----
Policy-id      :1
Description    :Default Security Queue policy

-----
Q      CIR      PIR      CBS      MBS      HiPrio
-----
1      1500      400000  10      5000000  10
2      3500      400000  40      5000000  10
-----
Associations
-----
MDA              :1/1 (Network Ingress)
MDA              :1/1 (Access Ingress)
MDA              :1/3 (Network Ingress)
MDA              :1/3 (Access Ingress)
MDA              :1/4 (Network Ingress)
MDA              :1/4 (Access Ingress)
MDA              :1/5 (Network Ingress)
MDA              :1/5 (Access Ingress)
MDA              :1/6 (Network Ingress)
MDA              :1/6 (Access Ingress)

-----
Security Queue Policy Id(2)
-----
Policy-id      :2
Description    :Description for Security Queue Policy id #2

-----
Q      CIR      PIR      CBS      MBS      HiPrio
-----
1      1500      400000  10      5000000  10
2      3500      400000  40      5000000  10
-----
Associations
-----
MDA              :1/2 (Access Ingress)

-----
Security Queue Policy Id(3)
-----
Policy-id      :3
Description    :Description for Security Queue Policy id #3
-----
```

```

Q      CIR      PIR      CBS      MBS      HiPrio
-----
1      1500     400000  10      5000000  10
2      3500     400000  40      5000000  10
-----
Associations
-----
MDA          :1/2 (Network Ingress)
=====
*A:7705custDoc:Sar18>show>qos#

```

Table 68: Security policy field descriptions

Label	Description
QoS Security Queue Policy	
Policy-id	The ID that uniquely identifies the security queue policy
Description	A text string that helps identify the security queue policy's context in the configuration file
Q	The security queue identifier, either 1 or 2
CIR	The committed information rate for the security queue
PIR	The peak information rate for the security queue
CBS	The committed buffer space for the security queue
MBS	The maximum burst size for the security queue
HiPrio	The percentage of the queue used exclusively by high-priority packets
Associations	
MDA	The adapter card slot number indicating the direction of traffic to which the security queue applies

12 List of acronyms

Table 69: Acronyms

Acronym	Expansion
2G	second-generation wireless telephone technology
3DES	triple DES (data encryption standard)
3G	third-generation mobile telephone technology
6VPE	IPv6 on virtual private edge router
7705 SAR	7705 Service Aggregation Router
7750 SR	7750 Service Router
8 PSK	eight phase shift keying
16 QAM	16-state quadrature amplitude modulation
32 QAM	32-state quadrature amplitude modulation
64 QAM	64-state quadrature amplitude modulation
128 QAM	128-state quadrature amplitude modulation
256 QAM	256-state quadrature amplitude modulation
ABR	area border router available bit rate
AC	alternating current attachment circuit
ACK	acknowledge
ACL	access control list
ACR	adaptive clock recovery
AD	auto-discovery
ADC	asymmetric delay control
ADM	add/drop multiplexer
ADP	automatic discovery protocol
AES	advanced encryption standard

Acronym	Expansion
AFI	authority and format identifier
AIGP	accumulated IGP
AIS	alarm indication signal
ALG	application level gateway
AMP	active multipath
AN	association number
ANSI	American National Standards Institute
Apipe	ATM VLL
APS	automatic protection switching
ARP	address resolution protocol
A/S	active/standby
AS	autonomous system
ASAP	any service, any port
ASBR	autonomous system boundary router
ASM	any-source multicast autonomous system message
ASN	autonomous system number
ATM	asynchronous transfer mode
ATM PVC	ATM permanent virtual circuit
AU	administrative unit
AUG	administrative unit group
B3ZS	bipolar with three-zero substitution
Batt A	battery A
B-bit	beginning bit (first packet of a fragment)
BBE	background block errors
Bc	committed burst size
Be	excess burst size
BECN	backward explicit congestion notification

Acronym	Expansion
Bellcore	Bell Communications Research
BFD	bidirectional forwarding detection
BGP	border gateway protocol
BGP-LS	border gateway protocol link state
BGP-LU	border gateway protocol labeled unicast
BITS	building integrated timing supply
BTCA	best timeTransmitter clock algorithm
BMU	broadcast, multicast, and unknown traffic Traffic that is not unicast. Any nature of multipoint traffic: <ul style="list-style-type: none"> • broadcast (that is, all 1s as the destination IP to represent all destinations within the subnet) • multicast (that is, traffic typically identified by the destination address, uses special destination address); for IP, the destination must be 224.0.0.0 to 239.255.255.255 • unknown (that is, the destination is typically a valid unicast address but the destination port/interface is not yet known; therefore, traffic needs to be forwarded to all destinations; unknown traffic is treated as broadcast)
BNM	bandwidth notification message
BOF	boot options file
BoS	bottom of stack
BPDU	bridge protocol data unit
BRAS	Broadband Remote Access Server
BSC	Base Station Controller
BSM	bootstrap message
BSR	bootstrap router
BSTA	Broadband Service Termination Architecture
BTS	base transceiver station
CA	certificate authority connectivity association
CAK	connectivity association key

Acronym	Expansion
CAS	channel associated signaling
CBN	common bonding networks
CBS	committed buffer space
CC	continuity check control channel
CCM	continuity check message
CCTV	closed-circuit television
CE	circuit emulation customer edge
CEM	circuit emulation
CES	circuit emulation services
CESoPSN	circuit emulation services over packet switched network
CFM	connectivity fault management
cHDLC	Cisco high-level data link control protocol
CIDR	classless inter-domain routing
CIR	committed information rate
CKN	connectivity association key name
CLI	command line interface
CLP	cell loss priority
CMP	certificate management protocol
C-multicast	customer multicast
CoS	class of service
CPE	customer premises equipment
Cpipe	circuit emulation (or TDM) VLL
CPM	Control and Processing Module (CPM is used instead of CSM when referring to CSM filtering to align with CLI syntax used with other SR products). CSM management ports are referred to as CPM management ports in the CLI.
CPROTO	C prototype

Acronym	Expansion
CPU	central processing unit
C/R	command/response
CRC	cyclic redundancy check
CRC-32	32-bit cyclic redundancy check
CRL	certificate revocation list
CRON	a time-based scheduling service (from chronos = time)
CRP	candidate RP
CSM	Control and Switching Module
CSNP	complete sequence number PDU
CSPF	constrained shortest path first
C-tag	customer VLAN tag
CV	connection verification customer VLAN (tag)
CW	control word
CWDM	coarse wavelength-division multiplexing
DA/FAN	distribution automation and field area network
DC	direct current
DC-C	DC return - common
DCE	data communications equipment
DC-I	DC return - isolated
DCO	digitally controlled oscillator
DCR	differential clock recovery
DDoS	distributed DoS
DE	discard eligibility
DER	distinguished encoding rules
DES	data encryption standard
DF	do not fragment designated forwarder

Acronym	Expansion
DH	Diffie-Hellman
DHB	decimal, hexadecimal, or binary
DHCP	dynamic host configuration protocol
DHCPv6	dynamic host configuration protocol for IPv6
DIS	designated intermediate system
DLCI	data link connection identifier
DLCMI	data link connection management interface
DM	delay measurement
DNS	domain name server
DNU	do not use
DoS	denial of service
dot1p	IEEE 802.1p bits, in Ethernet or VLAN ingress packet headers, used to map traffic to up to eight forwarding classes
dot1q	IEEE 802.1q encapsulation for Ethernet interfaces
DPD	dead peer detection
DPI	deep packet inspection
DPLL	digital phase locked loop
DR	designated router
DSA	digital signal algorithm
DSCP	differentiated services code point
DSL	digital subscriber line
DSLAM	digital subscriber line access multiplexer
DTE	data termination equipment
DU	downstream unsolicited
DUID	DHCP unique identifier
DUS	do not use for synchronization
DV	delay variation
DVMRP	distance vector multicast routing protocol

Acronym	Expansion
e911	enhanced 911 service
EAP	Extensible Authentication Protocol
EAPOL	EAP over LAN
E-bit	ending bit (last packet of a fragment)
E-BSR	elected BSR
ECMP	equal cost multipath
EE	end entity
EFM	Ethernet in the first mile
EGP	exterior gateway protocol
EIA/TIA-232	Electronic Industries Alliance/Telecommunications Industry Association Standard 232 (also known as RS-232)
EIR	excess information rate
EJBCA	Enterprise Java Beans Certificate Authority
E-LAN	Ethernet local area network
E-Line	Ethernet virtual private line
EL	entropy label
eLER	egress label edge router
ELI	entropy label indicator
E&M	ear and mouth earth and magneto exchange and multiplexer
eMBMS	evolved MBMS
EOP	end of packet
EPC	evolved packet core
EPD	early packet discard
Epip	Ethernet VLL
EPL	Ethernet private line
EPON	Ethernet Passive Optical Network

Acronym	Expansion
EPS	equipment protection switching
ERO	explicit route object
ES	Ethernet segment errored seconds
ESD	electrostatic discharge
ESI	Ethernet segment identifier
ESMC	Ethernet synchronization message channel
ESN	extended sequence number
ESP	encapsulating security payload
ESPI	encapsulating security payload identifier
ETE	end-to-end
ETH-BN	Ethernet bandwidth notification
ETH-CFM	Ethernet connectivity fault management (IEEE 802.1ag)
EVC	Ethernet virtual connection
EVDO	evolution - data optimized
EVI	EVPN instance
EVPL	Ethernet virtual private link
EVPN	Ethernet virtual private network
EXP bits	experimental bits (currently known as TC)
FC	forwarding class
FCS	frame check sequence
FD	frequency diversity
FDB	forwarding database
FDL	facilities data link
FEAC	far-end alarm and control
FEC	forwarding equivalence class
FECN	forward explicit congestion notification
FeGW	far-end gateway

Acronym	Expansion
FEP	front-end processor
FF	fixed filter
FFD	fast fault detection
FIB	forwarding information base
FIFO	first in, first out
FM	fault management
FNG	fault notification generator
FOM	figure of merit
Fpipe	frame relay VLL
FQDN	fully qualified domain name
FR	frame relay
FRG bit	fragmentation bit
FRR	fast reroute
FTN	FEC-to-NHLFE
FTP	file transfer protocol
FXO	foreign exchange office
FXS	foreign exchange subscriber
GFP	generic framing procedure
GigE	Gigabit Ethernet
GLONASS	Global Navigation Satellite System (Russia)
GNSS	global navigation satellite system (generic)
GPON	Gigabit Passive Optical Network
GPRS	general packet radio service
GPS	Global Positioning System
GRE	generic routing encapsulation
GRT	global routing table
GSM	Global System for Mobile Communications (2G)

Acronym	Expansion
GTP-U	GPRS tunneling protocol user plane
GW	gateway
HA	high availability
HCM	high capacity multiplexing
HDB3	high density bipolar of order 3
HDLC	high-level data link control protocol
HEC	header error control
HMAC	hash message authentication code
Hpipe	HDLC VLL
H-QoS	hierarchical quality of service
HSB	hot standby
HSDPA	high-speed downlink packet access
HSPA	high-speed packet access
H-VPLS	hierarchical virtual private line service
IANA	Internet Assigned Numbers Authority
IBN	isolated bonding networks
ICB	inter-chassis backup
ICK	integrity connection value key
ICMP	Internet control message protocol
ICMPv6	Internet control message protocol for IPv6
ICP	IMA control protocol cells
ICV	integrity connection value
IDS	intrusion detection system
IDU	indoor unit
IED	intelligent end device
IEEE	Institute of Electrical and Electronics Engineers
IEEE 1588v2	Institute of Electrical and Electronics Engineers standard 1588-2008

Acronym	Expansion
IES	Internet enhanced service
IETF	Internet Engineering Task Force
IGMP	Internet group management protocol
IGP	interior gateway protocol
IID	instance ID
IKE	Internet key exchange
iLER	ingress label edge router
ILM	incoming label map
IMA	inverse multiplexing over ATM
IMET-IR	inclusive multicast Ethernet tag—ingress replication
INVARP	inverse address resolution protocol
IOM	input/output module
IP	Internet protocol
IPCP	Internet protocol control protocol
IPIP	IP in IP
Ipipe	IP interworking VLL
I-PMSI	inclusive PMSI
IPoATM	IP over ATM
IPS	intrusion prevention system
IPSec	Internet protocol security
IR	ingress replication
IRB	integrated routing and bridging
ISA	integrated services adapter
ISAKMP	Internet security association and key management protocol
IS-IS	Intermediate System-to-Intermediate System
IS-IS-TE	IS-IS-traffic engineering (extensions)
ISO	International Organization for Standardization

Acronym	Expansion
IW	interworking
JP	join prune
KEK	key encryption key
KG	key group
LB	loopback
lbf-in	pound force inch
LBM	loopback message
LBO	line buildout
LBR	loopback reply
LCP	link control protocol
LDP	label distribution protocol
LER	label edge router
LFA	loop-free alternate
LFIB	label forwarding information base
LIB	label information base
LLDP	link layer discovery protocol
LLDPDU	link layer discovery protocol data unit
LLF	link loss forwarding
LLID	loopback location ID
LM	loss measurement
LMI	local management interface
LOS	line-of-sight loss of signal
LSA	link-state advertisement
LSDB	link-state database
LSP	label switched path link-state PDU (for IS-IS)
LSPA	LSP attributes

Acronym	Expansion
LSR	label switching router link-state request
LSU	link-state update
LT	linktrace
LTE	long term evolution line termination equipment
LTM	linktrace message
LTN	LSP ID to NHLFE
LTR	link trace reply
MA	maintenance association
MAC	media access control
MACsec	media access control security
MA-ID	maintenance association identifier
MBB	make-before-break
MBGP	multicast BGP multiprotocol BGP multiprotocol extensions for BGP
MBMS	multimedia broadcast multicast service
MBS	maximum buffer space maximum burst size media buffer space
MBSP	mobile backhaul service provider
MCAC	multicast connection admission control
MC-APS	multi-chassis automatic protection switching
MC-MLPPP	multiclass multilink point-to-point protocol
MCS	multicast server multi-chassis synchronization
MCT	MPT craft terminal
MD	maintenance domain

Acronym	Expansion
MD5	message digest version 5 (algorithm)
MDA	media dependent adapter
MDDDB	multidrop data bridge
MDL	maintenance data link
MDT	multicast distribution tree
ME	maintenance entity
MED	multi-exit discriminator
MEF	Metro Ethernet Forum
MEG	maintenance entity group
MEG-ID	maintenance entity group identifier
MEN	Metro Ethernet network
MEP	maintenance association endpoint
MFC	multi-field classification
MHD	multi-homed device
MHF	MIP half function
MHN	multi-homed network
MI	member identifier
MIB	management information base
MI-IS-IS	multi-instance IS-IS
MIR	minimum information rate
MKA	MACsec key agreement
MLD	multicast listener discovery
mLDP	multicast LDP
MLPPP	multilink point-to-point protocol
mLSP	multicast LSP
MoFRR	multicast-only fast reroute
MP	merge point multilink protocol

Acronym	Expansion
	multipoint
MP-BGP	multiprotocol border gateway protocol
MPLS	multiprotocol label switching
MPLSCP	multiprotocol label switching control protocol
MPP	MPT protection protocol
MPR	see Wavence
MPR-e	Microwave Packet Radio (standalone mode)
MPT-HC V2/9558HC	Microwave Packet Transport, High Capacity version 2
MPT-HLC	Microwave Packet Transport, High-Capacity Long-Haul Cubic (ANSI)
MPT-HQAM	Microwave Packet Transport, High Capacity (MPT-HC-QAM) or Extended Power (MPT-XP-QAM) with 512/1024 QAM
MPT-MC	Microwave Packet Transport, Medium Capacity
MPT-XP	Microwave Packet Transport, High Capacity (very high power version of MPT-HC V2/9558HC)
MRAI	minimum route advertisement interval
MRRU	maximum received reconstructed unit
MRU	maximum receive unit
MSDP	Multicast Source Discovery Protocol
MSDU	MAC Service Data Unit
MSK	master session key
MSO	multi-system operator
MS-PW	multi-segment pseudowire
MSS	maximum segment size Microwave Service Switch
MTIE	maximum time interval error
MTSO	mobile trunk switching office
MTU	maximum transmission unit multi-tenant unit

Acronym	Expansion
M-VPLS	management virtual private line service
MVPN	multicast VPN
MVR	multicast VPLS registration
MW	microwave
MWA	microwave awareness
N·m	newton meter
NAT	network address translation
NAT-T	network address translation traversal
NBMA	non-broadcast multiple access (network)
ND	neighbor discovery
NE	network element
NET	network entity title
NFM-P	Network Functions Manager - Packet (formerly 5620 SAM)
NGE	network group encryption
NG-MVPN	next generation MVPN
NH	next hop
NHLFE	next hop label forwarding entry
NHOP	next-hop
NLOS	non-line-of-sight
NLPID	network level protocol identifier
NLRI	network layer reachability information
NNHOP	next next-hop
NNI	network-to-network interface
Node B	similar to BTS but used in 3G networks — term is used in UMTS (3G systems) while BTS is used in GSM (2G systems)
NOC	network operations center
NPAT	network port address translation
NRC-F	Network Resource Controller - Flow

Acronym	Expansion
NRC-P	Network Resource Controller - Packet
NRC-T	Network Resource Controller - Transport
NRC-X	Network Resource Controller - Cross Domain
NSAP	network service access point
NSD	Network Services Director
NSP	native service processing Network Services Platform
NSSA	not-so-stubby area
NTP	Network Time Protocol
NTR	network timing reference
OADM	optical add/drop multiplexer
OAM	operations, administration, and maintenance
OAMPDU	OAM protocol data units
OC3	optical carrier level 3
OCSP	online certificate status protocol
ODU	outdoor unit
OIF	outgoing interface
OLT	optical line termination
OMC	optical management console
ONT	optical network terminal
OOB	out-of-band
OPX	off premises extension
ORF	outbound route filtering
OS	operating system
OSI	Open Systems Interconnection (reference model)
OSINLCP	OSI Network Layer Control Protocol
OSPF	open shortest path first
OSPF-TE	OSPF-traffic engineering (extensions)

Acronym	Expansion
OSS	operations support system
OSSP	organization specific slow protocol
OTP	one time password
OWAMP	one-way active measurement protocol
P2MP	point to multipoint
PADI	PPPoE active discovery initiation
PADR	PPPoE active discovery request
PAE	port authentication entities
PSB	path state block
PBO	packet byte offset
PBR	policy-based routing
PBX	private branch exchange
PCAP	packet capture
PCC	path computation client
PCE	path computation element
PCEP	Path Computation Element Communication Protocol
PCM	pulse code modulation
PCP	priority code point
PCR	proprietary clock recovery
PDU	power distribution unit protocol data units
PDV	packet delay variation
PDVT	packet delay variation tolerance
PE	provider edge router
PEAPv0	protected extensible authentication protocol version 0
PEM	privacy enhanced mail
PFoE	power feed over Ethernet
PFS	perfect forward secrecy

Acronym	Expansion
PHB	per-hop behavior
PHP	penultimate hop popping
PHY	physical layer
PIC	prefix independent convergence
PID	protocol ID
PIM SSM	protocol independent multicast—source-specific multicast
PIR	peak information rate
PKCS	public key cryptography standards
PKI	public key infrastructure
PLAR	private line automatic ringdown
PLCP	Physical Layer Convergence Protocol
PLR	point of local repair
PLSP	path LSP
PM	performance monitoring
PMSI	P-multicast service interface
P-multicast	provider multicast
PN	packet number
PoE	power over Ethernet
PoE+	power over Ethernet plus
POH	path overhead
POI	purge originator identification
PoP	point of presence
POS	packet over SONET
PPP	point-to-point protocol
PPPoE	point-to-point protocol over Ethernet
PPS	pulses per second
PRC	primary reference clock

Acronym	Expansion
PRS	primary reference source
PRTC	primary reference time clock
PSE	power sourcing equipment
PSK	pre-shared key
PSN	packet switched network
PSNP	partial sequence number PDU
PTA	PMSI tunnel attribute
PTM	packet transfer mode
PTP	performance transparency protocol Precision Time Protocol
PuTTY	an open-source terminal emulator, serial console, and network file transfer application
PVC	permanent virtual circuit
PVCC	permanent virtual channel connection
PW	pseudowire
PWE	pseudowire emulation
PWE3	pseudowire emulation edge-to-edge
Q.922	ITU-T Q-series Specification 922
QL	quality level
QoS	quality of service
QPSK	quadrature phase shift keying
RADIUS	Remote Authentication Dial In User Service
RAN	radio access network
RBS	robbed bit signaling
RD	route distinguisher
RDI	remote defect indication
RED	random early discard
RESV	reservation

Acronym	Expansion
RIB	routing information base
RIP	routing information protocol
RJ45	registered jack 45
RMON	remote network monitoring
RNC	radio network controller
RP	rendezvous point
RPF RTM	reverse path forwarding RTM
RPS	radio protection switching
RPT	rendezvous-point tree
RR	route reflector
RRO	record route object
RS-232	Recommended Standard 232 (also known as EIA/TIA-232)
RSA	Rivest, Shamir, and Adleman (authors of the RSA encryption algorithm)
RSHG	residential split horizon group
RSTP	rapid spanning tree protocol
RSVP-TE	resource reservation protocol - traffic engineering
RT	receive/transmit
RTC	route target constraint
RTM	routing table manager
RTN	battery return
RTP	real-time protocol
R&TTE	Radio and Telecommunications Terminal Equipment
RTU	remote terminal unit
RU	rack unit
r-VPLS	routed virtual private LAN service
SA	security association source-active

Acronym	Expansion
SAA	service assurance agent
SAFI	subsequent address family identifier
SAK	security association key
SAP	service access point
SAToP	structure-agnostic TDM over packet
SCADA	supervisory control and data acquisition
SC-APS	single-chassis automatic protection switching
SCI	secure channel identifier
SCP	secure copy
SCTP	Stream Control Transmission Protocol
SD	signal degrade space diversity
SDH	synchronous digital hierarchy
SDI	serial data interface
SDN	software defined network
SDP	service destination point
SE	shared explicit
SeGW	secure gateway
SES	severely errored seconds
SETS	synchronous equipment timing source
SF	signal fail
SFP	small form-factor pluggable (transceiver)
SFTP	SSH file transfer protocol
(S,G)	(source, group)
SGT	self-generated traffic
SHA-1	secure hash algorithm
SHG	split horizon group
SIR	sustained information rate

Acronym	Expansion
SL	short length
SLA	service-level agreement
SLARP	serial line address resolution protocol
SLID	subscriber location identifier of a GPON module
SLM	synthetic loss measurement
SNMP	Simple Network Management Protocol
SNPA	subnetwork point of attachment
SNR	signal to noise ratio
SNTP	simple network time protocol
SONET	synchronous optical networking
S-PE	switching provider edge router
SPF	shortest path first
SPI	security parameter index
S-PMSI	selective PMSI
SPT	shortest path tree
SR	Service Router (7750 SR) segment routing
SRGB	segment routing global block
SRLG	shared risk link group
SRP	stateful request parameter
SRRP	subscriber routed redundancy protocol
SR-ISIS	segment routing IS-IS
SR-OSPF	segment routing OSPF
SR-TE	segment routing traffic engineering
SSH	secure shell
SSM	source-specific multicast synchronization status messaging
SSU	system synchronization unit

Acronym	Expansion
S-tag	service VLAN tag
STM	synchronous transport module
STM1	synchronous transport module, level 1
STP	spanning tree protocol
STS	synchronous transport signal
SVC	switched virtual circuit
SVEC	synchronization vector
SYN	synchronize
TACACS+	Terminal Access Controller Access-Control System Plus
TC	traffic class (formerly known as EXP bits)
TCI	tag control information
TCP	transmission control protocol
TCP-AO	TCP Authentication Option
TDA	transmit diversity antenna
TDEV	time deviation
TDM	time division multiplexing
TE	traffic engineering
TEDB	traffic engineering database
TEID	tunnel endpoint identifier
TEP	tunnel endpoint
TFTP	trivial file transfer protocol
T-LDP	targeted LDP
TLS	transport layer security
TLV	type length value
TM	traffic management
ToD	time of day
ToS	type of service
T-PE	terminating provider edge router

Acronym	Expansion
TPID	tag protocol identifier
TPIF	IEEE C37.94 teleprotection interface
TPMR	two-port MAC relay
TPS	transmission protection switching
TSoP	transparent SDH/SONET over packet
TTL	time to live
TTLS	tunneled transport layer security
TTM	tunnel table manager
TU	tributary unit
TUG	tributary unit group
TWAMP	two-way active measurement protocol
U-APS	unidirectional automatic protection switching
UAS	unavailable seconds
UBR	unspecified bit rate
UDP	user datagram protocol
UFD	unidirectional forwarding detection
UMH	upstream multicast hop
UMTS	Universal Mobile Telecommunications System (3G)
UNI	user-to-network interface
uRPF	unicast reverse path forwarding
V.11	ITU-T V-series Recommendation 11
V.24	ITU-T V-series Recommendation 24
V.35	ITU-T V-series Recommendation 35
VC	virtual circuit
VCB	voice conference bridge
VCC	virtual channel connection
VCCV	virtual circuit connectivity verification

Acronym	Expansion
VCI	virtual circuit identifier
VID	VLAN ID
VLAN	virtual LAN
VLL	virtual leased line
VM	virtual machine
VoIP	voice over IP
Vp	peak voltage
VP	virtual path
VPC	virtual path connection
VPI	virtual path identifier
VPLS	virtual private LAN service
VPN	virtual private network
VPRN	virtual private routed network
VPWS	virtual private wire service
VRF	virtual routing and forwarding table
VRRP	virtual router redundancy protocol
V-SAP	virtual service access point
VSE	vendor-specific extension
VSI	virtual switch instance
VSO	vendor-specific option
VT	virtual trunk virtual tributary
VTG	virtual tributary group
Wavence	formerly 9500 MPR (Microwave Packet Radio)
WCDMA	wideband code division multiple access (transmission protocol used in UMTS networks)
WRED	weighted random early discard
WTR	wait to restore

Acronym	Expansion
X.21	ITU-T X-series Recommendation 21
XOR	exclusive-OR
XRO	exclude route object

13 Supported standards and protocols

This chapter lists the 7705 SAR compliance with security and telecom standards, the protocols supported, and proprietary MIBs.

13.1 Telecom standards

ANSI/TIA/EIA-232-C—Interface Between Data Terminal Equipment and Data Circuit-Terminating Equipment Employing Serial Binary Data Interchange

IEEE 802.1AB-2016—IEEE Standard for Local and metropolitan area networks - Station and Media Access Control Connectivity Discovery

IEEE 802.1ad—IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks

IEEE Std 802.1AE-2006 Media Access Control (MAC) Security

IEEE Std 802.1AEbw-2013—Media Access Control (MAC) Security Amendment 2: Extended Packet Numbering

IEEE 802.1ag—Service Layer OAM

IEEE 802.1p/q—VLAN Tagging

IEEE 802.1x-2010—IEEE Standard for Local and Metropolitan Area Networks—Port-based Network Access Control

IEEE 802.3—10BaseT

IEEE 802.3ab—1000BaseT

IEEE 802.3ah—Ethernet OAM

IEEE 802.3u—100BaseTX

IEEE 802.3x —Flow Control

IEEE 802.3z—1000BaseSX/LX

IEEE 802.3-2008—Revised base standard

IEEE 802.1AX-2008—Link Aggregation Task Force (transferred from IEEE 802.3ad)

IEEE C37.94-2017—N Times 64 Kilobit Per Second Optical Fiber Interfaces Between Teleprotection and Multiplexer Equipment

ITU-T G.704—Synchronous frame structures used at 1544, 6312, 2048, 8448 and 44 736 kbit/s hierarchical levels

ITU-T G.707—Network node interface for the Synchronous Digital Hierarchy (SDH)

ITU-T G.826—End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections

ITU-T G.8032 — Ethernet Ring Protection Switching

ITU-T G.984.1—Gigabit-capable passive optical networks (GPON): general characteristics

ITU-T Y.1564—Ethernet service activation test methodology

ITU-T Y.1731—OAM functions and mechanisms for Ethernet-based networks

13.2 Protocol support

13.2.1 ATM

AF-PHY-0086.001—Inverse Multiplexing for ATM (IMA)

af-tm-0121.000—Traffic Management Specification Version 4.1, March 1999

GR-1113-CORE—Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

GR-1248-CORE—Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996

ITU-T Recommendation I.432.1—B-ISDN user-network interface - Physical layer specification: General characteristics

ITU-T Recommendation I.610—B-ISDN Operation and Maintenance Principles and Functions version 11/95

RFC 2514—Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management, February 1999

RFC 2515—Definition of Managed Objects for ATM Management, February 1999

RFC 2684—Multiprotocol Encapsulation over ATM Adaptation Layer 5

13.2.2 BFD

RFC 7130—Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces

RFC 7881—Seamless Bidirectional Forwarding Detection (S-BFD) for IPv4, IPv6, and MPLS

draft-ietf-bfd-mib-00.txt—Bidirectional Forwarding Detection Management Information Base

draft-ietf-bfd-base-o5.txt—Bidirectional Forwarding Detection

draft-ietf-bfd-v4v6-1hop-06.txt—BFD IPv4 and IPv6 (Single Hop)

draft-ietf-bfd-multihop-06.txt—BFD for Multi-hop Paths

13.2.3 BGP

RFC 1397—BGP Default Route Advertisement

RFC 1997—BGP Communities Attribute

RFC 2385—Protection of BGP Sessions via the TCP MD5 Signature Option

RFC 2439—BGP Route Flap Dampening

RFC 2545—Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing

RFC 2918—Route Refresh Capability for BGP-4
RFC 3107—Carrying Label Information in BGP-4
RFC 3392—Capabilities Advertisement with BGP-4
RFC 4271—BGP-4 (previously RFC 1771)
RFC 4360—BGP Extended Communities Attribute
RFC 4364—BGP/MPLS IP Virtual Private Networks (VPNs) (previously RFC 2574bis BGP/MPLS VPNs)
RFC 4456—BGP Route Reflection: Alternative to Full-mesh IBGP (previously RFC 1966 and RFC 2796)
RFC 4486—Subcodes for BGP Cease Notification Message
RFC 4684—Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)
RFC 4724—Graceful Restart Mechanism for BGP - GR Helper
RFC 4760—Multi-protocol Extensions for BGP (previously RFC 2858)
RFC 4893—BGP Support for Four-octet AS Number Space
RFC 4798—Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)
RFC 5549—Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop
RFC 5925—The TCP Authentication Option
RFC 5926—Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)
RFC 6513—Multicast in MPLS/BGP IP VPNs
RFC 6514—BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs
RFC 7311—The Accumulated IGP Metric Attribute for BGP
RFC 7606—Revised Error Handling for BGP UPDATE Messages
draft-ietf-idr-add-paths-04.txt—Advertisement of Multiple Paths in BGP
draft-ietf-idr-add-paths-guidelines-00.txt—Best Practices for Advertisement of Multiple Paths in BGP
draft-weis-esp-group-counter-cipher-00.txt—Using Counter Modes with Encapsulating Security Payload (ESP) and Authentication Header (AH) to Protect Group Traffic

13.2.4 DHCP/DHCPv6

RFC 1534—Interoperation between DHCP and BOOTP
RFC 2131—Dynamic Host Configuration Protocol (REV)
RFC 2132—DHCP Options and BOOTP Vendor Extensions
RFC 3046—DHCP Relay Agent Information Option (Option 82)
RFC 3315—Dynamic Host Configuration Protocol for IPv6
RFC 3736—Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6

13.2.5 Differentiated services

RFC 2474—Definition of the DS Field in the IPv4 and IPv6 Headers

RFC 2597—Assured Forwarding PHB Group
RFC 2598—An Expedited Forwarding PHB
RFC 3140—Per-Hop Behavior Identification Codes

13.2.6 Digital data network management

V.35
RS-232 (also known as EIA/TIA-232)
X.21

13.2.7 ECMP

RFC 2992—Analysis of an Equal-Cost Multi-Path Algorithm

13.2.8 Ethernet VPN (EVPN)

RFC 7432—BGP MPLS-Based Ethernet VPN
draft-ietf-bess-evpn-vpws—Virtual Private Wire Service support in Ethernet VPN
draft-ietf-rabadan-bess-evpn-pref-pdf—Preference-based EVPN DF Election

13.2.9 Frame relay

ANSI T1.617 Annex D—Signalling Specification For Frame Relay Bearer Service
ITU-T Q.922 Annex A—Digital Subscriber Signalling System No. 1 (DSS1) data link layer - ISDN data link layer specification for frame mode bearer services
FRF.1.2—PVC User-to-Network Interface (UNI) Implementation Agreement
RFC 2427—Multiprotocol Interconnect over Frame Relay

13.2.10 GRE

RFC 2784—Generic Routing Encapsulation (GRE)

13.2.11 Internet protocol (IP) – version 4

RFC 768—User Datagram Protocol
RFC 791—Internet Protocol
RFC 792—Internet Control Message Protocol
RFC 793—Transmission Control Protocol
RFC 826—Ethernet Address Resolution Protocol

RFC 854—Telnet Protocol Specification
RFC 1350—The TFTP Protocol (Rev. 2)
RFC 1812—Requirements for IPv4 Routers
RFC 3021—Using 31-Bit Prefixes on IPv4 Point-to-Point Links

13.2.12 Internet protocol (IP) – version 6

RFC 2460—Internet Protocol, Version 6 (IPv6) Specification
RFC 2462—IPv6 Stateless Address Autoconfiguration
RFC 2464—Transmission of IPv6 Packets over Ethernet Networks
RFC 3587—IPv6 Global Unicast Address Format
RFC 3595—Textual Conventions for IPv6 Flow Label
RFC 4007—IPv6 Scoped Address Architecture
RFC 4193—Unique Local IPv6 Unicast Addresses
RFC 4291—IPv6 Addressing Architecture
RFC 4443—Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 Specification
RFC 4649—DHCPv6 Relay Agent Remote-ID Option
RFC 4861—Neighbor Discovery for IP version 6 (IPv6)
RFC 5095—Deprecation of Type 0 Routing Headers in IPv6
RFC 5952—A Recommendation for IPv6 Address Text Representation

13.2.13 IPSec

ITU-T X.690 (2002)—ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)
PKCS #12 Personal Information Exchange Syntax Standard
RFC 2315—PKCS #7: Cryptographic Message Syntax
RFC 2409—The Internet Key Exchange (IKE)
RFC 2986—PKCS #10: Certification Request Syntax Specification
RFC 3706—A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers
RFC 3947—Negotiation of NAT-Traversal in the IKE
RFC 3948—UDP Encapsulation of IPsec ESP Packets
RFC 4301—Security Architecture for the Internet Protocol
RFC 4303—IP Encapsulating Security Payload (ESP)
RFC 4210—Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)
RFC 4211—Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)
RFC 4945—The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2, and PKIX

RFC 5280—Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile

RFC 5996—Internet Key Exchange Protocol Version 2 (IKEv2)

RFC 7383—Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation

13.2.14 IS-IS

RFC 1142—OSI IS-IS Intra-domain Routing Protocol (ISO 10589)

RFC 1195—Use of OSI IS-IS for routing in TCP/IP & dual environments

RFC 2763—Dynamic Hostname Exchange for IS-IS

RFC 2966—Domain-wide Prefix Distribution with Two-Level IS-IS

RFC 2973—IS-IS Mesh Groups

RFC 3373—Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies

RFC 3567—Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication

RFC 3719—Recommendations for Interoperable Networks using IS-IS

RFC 3784—Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)

RFC 3787—Recommendations for Interoperable IP Networks

RFC 4205 for Shared Risk Link Group (SRLG) TLV

RFC 4971—Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information

RFC 5120—M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)

RFC 5304—IS-IS Cryptographic Authentication

RFC 5305—IS-IS Extensions for Traffic Engineering

RFC 5307—IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)

RFC 5308—Routing IPv6 with IS-IS

RFC 5309—Point-to-Point Operation over LAN in Link State Routing Protocols

RFC 5310—IS-IS Generic Cryptographic Authentication

RFC 6232—Purge Originator Identification TLV for IS-IS

13.2.15 LDP

RFC 5036—LDP Specification

RFC 5283—LDP Extension for Inter-Area Label Switched Paths

RFC 5350—IANA Considerations for the IPv4 and IPv6 Router Alert Options

RFC 5443—LDP IGP Synchronization

RFC 5561—LDP Capabilities

RFC 5926—Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)

RFC 6388—Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

RFC 6512—Using Multipoint LDP When the Backbone Has No Route to the Root

RFC 6829—Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6

RFC 7552—Updates to LDP for IPv6

draft-ietf-mpls-ldp-ip-pw-capability—Controlling State Advertisements Of Non-negotiated LDP Applications

draft-ietf-mpls-oam-ipv6-rao—IPv6 Router Alert Option for MPLS OAM

draft-pdutta-mpls-ldp-adj-capability-00—LDP Adjacency Capabilities

draft-pdutta-mpls-ldp-v2-00—LDP Version 2

draft-pdutta-mpls-mldp-up-redundancy-00.txt—Upstream LSR Redundancy for Multi-point LDP Tunnels

draft-weis-esp-group-counter-cipher-00—Using Counter Modes with Encapsulating Security Payload (ESP) and Authentication Header (AH) to Protect Group Traffic

13.2.16 LDP and IP FRR

RFC 5286—Basic Specification for IP Fast Reroute: Loop-Free Alternates

RFC 7490—Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)

13.2.17 MPLS

RFC 3031—MPLS Architecture

RFC 3032—MPLS Label Stack Encoding

RFC 3815—Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)

RFC 5440—Path Computation Element (PCE) Communication Protocol (PCEP)

RFC 6790—The Use of Entropy Labels in MPLS Forwarding

RFC 8253—PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)

RFC 8697—Path Computation Element Communication Protocol (PCEP) Extensions for Establishing Relationships between Sets of Label Switched Paths (LSPs)

RFC 8745—Path Computation Element Communication Protocol (PCEP) Extensions for Associating Working and Protection Label Switched Paths (LSPs) with Stateful PCE

RFC 8800—Path Computation Element Communication Protocol (PCEP) Extension for Label Switched Path (LSP) Diversity Constraint Signaling

draft-dhody-pce-pceps-tls13-02—Updates for PCEPS

draft-ietf-pce-stateful-pce—PCEP Extensions for Stateful PCE

draft-ietf-pce-segment-routing—PCEP Extensions for Segment Routing

draft-alvarez-pce-path-profiles—PCE Path Profiles

13.2.18 MPLS – OAM

RFC 6424— Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels

RFC 8029—Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

13.2.19 Multicast

RFC 3956—Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address

RFC 3973—Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)

RFC 4610—Anycast-RP Using Protocol Independent Multicast (PIM), which is similar to RFC 3446—Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)

RFC 6514—BGP Encodings and Procedures for Multicast in MPLS/IP VPNs

RFC 6826—Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

cisco-ipmulticast/pim-autorp-spec—Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast, which is similar to RFC 5059—Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)

draft-ietf-l2vpn-vpls-pim-snooping-07—Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)

draft-ietf-mboned-msdp-deploy-nn.txt—Multicast Source Discovery Protocol (MSDP) Deployment Scenarios

13.2.20 Network management

IANA-IFType-MIB

ITU-T X.721—Information technology- OSI-Structure of Management Information

ITU-T X.734—Information technology- OSI-Systems Management: Event Report Management Function

M.3100/3120—Equipment and Connection Models

RFC 1157—SNMPv1

RFC 1850—OSPF-MIB

RFC 1907—SNMPv2-MIB

RFC 2011—IP-MIB

RFC 2012—TCP-MIB

RFC 2013—UDP-MIB

RFC 2030—Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI

RFC 2096—IP-FORWARD-MIB

RFC 2138—RADIUS

RFC 2206—RSVP-MIB

RFC 2571—SNMP-FRAMEWORKMIB
RFC 2572—SNMP-MPD-MIB
RFC 2573—SNMP-TARGET-&-NOTIFICATION-MIB
RFC 2574—SNMP-USER-BASED-SMMIB
RFC 2575—SNMP-VIEW-BASED ACM-MIB
RFC 2576—SNMP-COMMUNITY-MIB
RFC 2588—SONET-MIB
RFC 2665—EtherLike-MIB
RFC 2819—RMON-MIB
RFC 2863—IF-MIB
RFC 2864—INVERTED-STACK-MIB
RFC 3014—NOTIFICATION-LOG MIB
RFC 3164—The BSD Syslog Protocol
RFC 3273—HCRMON-MIB
RFC 3411—An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks
RFC 3412—Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
RFC 3413—Simple Network Management Protocol (SNMP) Applications
RFC 3414—User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
RFC 3418—SNMP MIB
RFC 3954—Cisco Systems NetFlow Services Export Version 9
RFC 5101—Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information
RFC 5102—Information Model for IP Flow Information Export
draft-ietf-disman-alarm-mib-04.txt
draft-ietf-mpls-ldp-mib-07.txt
draft-ietf-ospf-mib-update-04.txt
draft-ietf-mpls-lsr-mib-06.txt
draft-ietf-mpls-te-mib-04.txt
TMF 509/613—Network Connectivity Model

13.2.21 OSPF

RFC 1765—OSPF Database Overflow
RFC 2328—OSPF Version 2
RFC 2370—Opaque LSA Support
RFC 2740—OSPF for IPv6

RFC 3101—OSPF NSSA Option
RFC 3137—OSPF Stub Router Advertisement
RFC 3509—Alternative Implementations of OSPF Area Border Routers
RFC 3623—Graceful OSPF Restart (support for Helper mode)
RFC 3630—Traffic Engineering (TE) Extensions to OSPF
RFC 4203 for Shared Risk Link Group (SRLG) sub-TLV
RFC 4577—OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) (support for basic OSPF at PE-CE links)
RFC 4915—Multi-Topology (MT) Routing in OSPF
RFC 4970—Extensions to OSPF for Advertising Optional Router Capabilities
RFC 5185—OSPF Multi-Area Adjacency

13.2.22 OSPFv3

RFC 4552—Authentication/Confidentiality for OSPFv3

13.2.23 PPP

RFC 1332—PPP Internet Protocol Control Protocol (IPCP)
RFC 1570—PPP LCP Extensions
RFC 1619—PPP over SONET/SDH
RFC 1661—The Point-to-Point Protocol (PPP)
RFC 1662—PPP in HDLC-like Framing
RFC 1989—PPP Link Quality Monitoring
RFC 1990—The PPP Multilink Protocol (MP)
RFC 2686—The Multi-Class Extension to Multi-Link PPP

13.2.24 Pseudowires

Metro Ethernet Forum—Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks
RFC 3550—RTP: A Transport Protocol for Real-Time Applications
RFC 3985—Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture
RFC 4385—Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN
RFC 4446—IANA Allocation for PWE3
RFC 4447—Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)
RFC 4448—Encapsulation Methods for Transport of Ethernet over MPLS Networks
RFC 4553—Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)

RFC 4717—Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks

RFC 4618—Encapsulation Methods for Transport of PPP/High-Level Data Link Control (HDLC) over MPLS Networks

RFC 4619—Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks

RFC 4816—Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service

RFC 5085—Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires

RFC 5086—Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)

draft-ietf-pwe3-redundancy-02.txt—Pseudowire (PW) Redundancy

13.2.25 RIP

RFC 1058—Routing Information Protocol

RFC 2453—RIP Version 2

13.2.26 RADIUS

RFC 2865—Remote Authentication Dial In User Service

RFC 2866—RADIUS Accounting

RFC 6613—RADIUS over TCP

RFC 6614—Transport Layer Security (TLS) Encryption for RADIUS

13.2.27 RSVP-TE and FRR

RFC 2430—A Provider Architecture for DiffServ & TE

RFC 2702—Requirements for Traffic Engineering over MPLS

RFC 2747—RSVP Cryptographic Authentication

RFC 2961—RSVP Refresh Overhead Reduction Extensions

RFC 3097—RSVP Cryptographic Authentication - Updated Message Type Value

RFC 3209—Extensions to RSVP for LSP Tunnels

RFC 3210—Applicability Statement for Extensions to RSVP for LSP Tunnels

RFC 3477—Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)

RFC 4090—Fast Reroute Extensions to RSVP-TE for LSP Tunnels

13.2.28 Segment routing (SR)

draft-francois-rtgwg-segment-routing-ti-lfa-04—Topology Independent Fast Reroute using Segment Routing

draft-gredler-idr-bgp-ls-segment-routing-ext-03—BGP Link-State extensions for Segment Routing

draft-ietf-isis-segment-routing-extensions-04—IS-IS Extensions for Segment Routing

draft-ietf-mpls-spring-lsp-ping-02—Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane

draft-ietf-ospf-segment-routing-extensions-04—OSPF Extensions for Segment Routing

draft-ietf-spring-segment-routing-15—Segment Routing Architecture

13.2.29 SONET/SDH

GR-253-CORE—SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

ITU-T Recommendation G.841—Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

13.2.30 SSH

RFC 4253—The Secure Shell (SSH) Transport Layer Protocol

draft-ietf-secsh-architecture.txt—SSH Protocol Architecture

draft-ietf-secsh-userauth.txt—SSH Authentication Protocol

draft-ietf-secsh-connection.txt—SSH Connection Protocol

draft-ietf-secsh-newmodes.txt—SSH Transport Layer Encryption Modes

draft-ietf-secsh-filexfer-13.txt—SSH File Transfer Protocol

13.2.31 Synchronization

G.781—Synchronization layer functions, 2001/09/17

G.803—Architecture of transport networks based on the synchronous digital hierarchy (SDH)

G.813—Timing characteristics of SDH equipment slave clocks (SEC)

G.823—The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy, 2003/03/16

G.824—The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy, 2003/03/16

G.8261—Timing and synchronization aspects in packet networks

G.8262—Timing characteristics of synchronous Ethernet equipment slave clock

GR 1244 CORE—Clocks for the Synchronized Network: Common Generic Criteria

IEC/IEEE 61850-9-3—Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation

IEEE C37.238-2017 - IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications

IEEE Std 1588-2008—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

IEEE Std 1588-2008—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Annex E – Transport of PTP over User Datagram Protocol over Internet Protocol Version 6

IEEE Std 1588-2019—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Annex J

ITU-T G.8264—Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008

ITU-T G.8265.1—Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for frequency synchronization, issued 10/2010

ITU-T G.8275.1—Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014

ITU-T G.8275.2—Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for time/phase synchronization with partial timing support from the network, issued 06/2016

RFC 5905—Network Time Protocol Version 4: Protocol and Algorithms Specification

RFC 8573—Message Authentication Code for the Network Time Protocol

13.2.32 TACACS+

IETF draft-grant-tacacs-02.txt—The TACACS+ Protocol

13.2.33 TLS

RFC 5246—The Transport Layer Security (TLS) Protocol Version 1.2

RFC 5425—Transport Layer Security (TLS) Transport Mapping for Syslog

RFC 5922—Domain Certificates in the Session Initiation Protocol (SIP)

RFC 6460—Suite B Profile for Transport Layer Security (TLS)

RFC 8446—The Transport Layer Security (TLS) Protocol Version 1.3

13.2.34 TWAMP

RFC 5357—A Two-Way Active Measurement Protocol (TWAMP)

13.2.35 VPLS

RFC 4762—Virtual Private LAN Services Using LDP

13.2.36 VRRP

RFC 2787—Definitions of Managed Objects for the Virtual Router Redundancy Protocol

RFC 3768 Virtual Router Redundancy Protocol

RFC 5798 Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6

13.3 Proprietary MIBs

TIMETRA-ATM-MIB.mib

TIMETRA-CAPABILITY-7705-V1.mib

TIMETRA-CHASSIS-MIB.mib

TIMETRA-CLEAR-MIB.mib

TIMETRA-FILTER-MIB.mib

TIMETRA-GLOBAL-MIB.mib

TIMETRA-LAG-MIB.mib

TIMETRA-LDP-MIB.mib

TIMETRA-LOG-MIB.mib

TIMETRA-MPLS-MIB.mib

TIMETRA-OAM-TEST-MIB.mib

TIMETRA-PORT-MIB.mib

TIMETRA-PPP-MIB.mib

TIMETRA-QOS-MIB.mib

TIMETRA-ROUTE-POLICY-MIB.mib

TIMETRA-RSVP-MIB.mib

TIMETRA-SAP-MIB.mib

TIMETRA-SDP-MIB.mib

TIMETRA-SECURITY-MIB.mib

TIMETRA-SERV-MIB.mib

TIMETRA-SYSTEM-MIB.mib

TIMETRA-TC-MIB.mib

TIMETRA-VRRP-MIB.mib

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)